Netra™ T2000 Server Administration Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 819-5837-10
September 2006, Revision A

Submit comments about this document at: http://www.sun.com/hwdocs/feedback

Contents

Figures

Tables

Preface

The *Netra T2000 Server Administration Guide* provides information and detailed procedures that enable administration and management of the Netra™ T2000 server. This document is written for technicians, system administrators, authorized service providers (ASPs), and users who have advanced experience administrating server systems.

How This Document Is Organized

Chapter 1 explains how to access the system console to enable remote management and administration.

Chapter 2 describes using Advanced Lights Out Manager (ALOM) for remote administration of your server.

Chapter 3 describes the function, methods of obtaining, and configuration of the OpenBoot™ PROM.

Chapter 4 describes status indicators and basic tasks that might be done as the course of system administration.

Chapter 5 provides important information about securing the system.

Chapter 6 describes redundant array of independent disks (RAID) concepts.

Appendix A gives information on the watchdog timer application mode on the server.

Appendix B provides a sample program that illustrates how to get or set the status of the alarms.

Using UNIX Commands

This document might not contain information about basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. Refer to the following for this information:

- Software documentation that you received with your system
- Solaris™ Operating System documentation, which is at:

  http://docs.sun.com

Shell Prompts

Shell	Prompt
C shell	*machine-name*%
C shell superuser	*machine-name*#
Bourne shell and Korn shell	$
Bourne shell and Korn shell superuser	#

Typographic Conventions

Typeface*	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your .login file. Use ls -a to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% **su** Password:
AaBbCc123	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be superuser to do this. To delete a file, type rm *filename*.

* The settings on your browser might differ from these settings.

Related Documentation

The documents listed as online are available at:

http://www.sun.com/products-n-solutions/hardware/docs/

Application	Title	Part Number	Format	Location
Installation	*Netra T2000 Server Installation Guide*	819-5838	PDF	Online
Updates	*Netra T2000 Server Product Notes*	819-5840	PDF	Online
Service	*Netra T2000 Server Service Manual*	819-5841	PDF	Online
Planning	*Netra T2000 Server Site Planning Notes*	819-5842	PDF	Online
Compliance	*Netra T2000 Server Safety and Compliance Guide*	819-5843	PDF	Online
Documentation	*Netra T2000 Server Getting Started Guide*	819-5844	Printed PDF	Shipping kit Online
Reference	*ALOM CMT 1.2 Guide*	819-3250	PDF	Online

Documentation, Support, and Training

Sun Function	URL
Documentation	http://www.sun.com/documentation/
Support	http://www.sun.com/support/
Training	http://www.sun.com/training/

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

http://www.sun.com/hwdocs/feedback

Please include the title and part number of your document with your feedback:

Netra T2000 Server Administration Guide, part number 819-5837-10

CHAPTER 1

System Console

This chapter explains how to access the system console to enable remote management and administration. This chapter includes the following topics:

Communicating With the System Console

The administrator needs a way to interact with the server at a low level, to configure the very basic I/O and server boot behavior. The system console enables the administrator to accomplish these tasks, using special commands. Additionally, the system console displays information, status, and error messages generated by firmware during server startup and operation. Once the operating system is booted, the system console displays Solaris system messages and accepts Solaris commands.

The server has two I/O ports dedicated to the system console:

- SC SERIAL MGT
- SC NET MGT

Serial Management Port

The serial management port (SC SERIAL MGT) is the default connection to the system console. Though a serial connection, this port uses an RJ-45 connector. Communicating with the system controller through this port requires the following serial parameters:

- 9600 baud

- 8 bits
- No parity
- 1 stop bit
- No handshaking

Serial devices that can communicate with the serial management port are:

- Terminal server
- TIP line connected to another Sun™ computer
- Alphanumeric terminal or similar device

Because it is a serial connection, there is communication between only two devices. This constraint limits access and provides for a more secure link between the administrator and the server.

The serial management port is not a general-purpose serial port. It is dedicated to the system controller. If you want to use a serial peripheral, connect it to the standard 9-pin serial port on the back panel of the server. The Solaris OS sees this port as TTYA and it is labeled as such.

Establishing Communication With the Serial Management Port

▼ To Access The System Console Through a Terminal Server

1. **Complete the physical connection from the serial management port to your terminal server.**

   The serial management port on the server is a data terminal equipment (DTE) port. Verify that the serial port pinouts of the server match those of the terminal server you plan to use.

   - If the pinout of the server serial management port corresponds with the pinout of the RJ-45 port on the terminal server, you have two connection options:
     - Connect a serial interface breakout cable directly to the server.
     - Connect a serial interface breakout cable to a patch panel and use the straight-through patch cable (supplied by Sun) to connect the patch panel to the server.
   - If the pinout of the server serial management port *do not* correspond with the pinout of the RJ-45 port on the terminal server, you need to make a crossover cable. TABLE 1-1 shows the pinout of the crossover cable.

TABLE 1-1 Pin Crossovers for Connecting to a Typical Terminal Server

Server Serial Port (RJ-45 Connector) Pin	Terminal Server Serial Port Pin
Pin 1 (RTS)	Pin 1 (CTS)
Pin 2 (DTR)	Pin 2 (DSR)
Pin 3 (TXD)	Pin 3 (RXD)
Pin 4 (Signal Ground)	Pin 4 (Signal Ground)
Pin 5 (Signal Ground)	Pin 5 (Signal Ground)
Pin 6 (RXD)	Pin 6 (TXD)
Pin 7 (DSR /DCD)	Pin 7 (DTR)
Pin 8 (CTS)	Pin 8 (RTS)

2. **Open a terminal session on the connecting device, and type:**

```
% telnet IP-address-of-terminal-server  port-number
```

For example, for a server connected to port 10000 on a terminal server whose IP address is 192.20.30.10, you would type:

```
% telnet 192.20.30.10 10000
```

▼ To Access the System Console Through the TIP Connection

1. **Connect the RJ-45 serial cable and, if required, the DB-9 or DB-25 adapter provided.**

   The cable and adapter connect between another Sun system's serial port (typically TTYB) and the serial management port on the back panel of the server.

2. **Ensure that the `/etc/remote` file on the Sun system contains an entry for `hardwire`.**

   See TABLE 1-2.

TABLE 1-2 Entries for `hardwire` In the `/etc/remote` File

Serial port	Entry for `hardwire`
`ttya`	`hardwire:\` `:dv=/dev/term/a:br#9600:el=^C^S^Q^U^D:ie=%$:oe=^D:`
`ttyb`	`hardwire:\` `:dv=/dev/term/b:br#9600:el=^C^S^Q^U^D:ie=%$:oe=^D:`

3. **In a terminal window on the Sun system, type:**

```
% tip hardwire
```

   The Sun system responds by displaying:

```
connected
```

   The server and the Sun system are now communicating.

▼ To Access the System Console Through an Alphanumeric Terminal

1. **Attach one end of the serial cable to the alphanumeric terminal's serial port.**

   Use a null modem serial cable, or an RJ-45 serial cable and null modem adapter. Connect this cable to the terminal's serial port connector.

2. **Attach the opposite end of the serial cable to the serial management port on the server.**

3. **Power on the alphanumeric terminal.**

4. **Set the alphanumeric terminal to receive:**

   - 9600 baud
   - 8 bits
   - No parity
   - 1 stop bit
   - No handshake protocol

   Refer to the documentation accompanying your terminal for information about how to operate and configure the terminal.

Network Management Port

The network management port (SC NET MGT) permits communication with the system controller, through your existing Ethernet network. The network management port is a 10/100BASE-T port with a unique IP address, separate from the server IP address. Like the serial management port, the network management port is dedicated to the system controller. Unlike the serial management port, up to eight sessions of the system controller can exist concurrently. As such, strict control of system controller access is required.

Before you can use the network management port, its unique IP address must be assigned using the serial management port. You can either assign a static IP address or configure the system controller to find a dynamic IP address, using DHCP.

Note – Data centers frequently devote a separate subnet to system management. If your data center has such a configuration, connect the network management port to this subnet.

▼ To Activate the Network Management Port

1. **Connect an Ethernet cable to the network management port.**
2. **Log in to the system controller through the serial management port.**

   See "Establishing Communication With the Serial Management Port" on page 2.
3. **Type one of the following commands:**
   - **If your network uses static IP addresses, type:**

```
sc> setsc if_network true
sc> setsc netsc_ipaddr ip-address
sc> setsc netsc_ipnetmask ip-address
sc> setsc netsc_ipgateway ip-address
```

   - **If your network uses DHCP, type:**

```
sc> setsc netsc_dhcp true
```

4. **Reset the system controller so that the new settings take affect:**

```
sc> resetsc
```

5. **After the system controller resets, use the `shownetwork` command to verify network settings:**

```
sc> shownetwork
```

6. **Exit the system controller session.**

```
sc> console
```

To connect through the network management port, use the `telnet` command to the IP address you specified in Step 3 of "To Activate the Network Management Port" on page 5.

Switching Between the Consoles

The system controller (SC) console connection provides access to the ALOM shell, the Solaris OS, and the OpenBoot PROM.

This section describes the procedures to navigate between the following:

- ALOM prompt (`sc>`)
- Solaris OS prompt (`#`)
- OpenBoot PROM prompt (`ok`)

These procedures are summarized in FIGURE 1-1

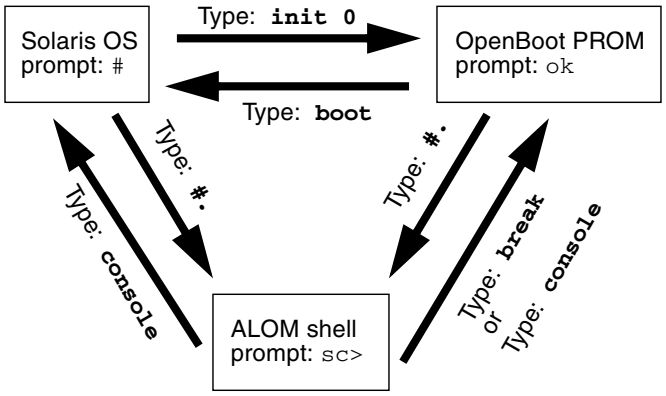

FIGURE 1-1 Navigation Between Consoles

ALOM sc> Prompt

The ALOM runs independently of the server and regardless of the server power state. When you connect the server to source power, the ALOM immediately starts up and begins monitoring the system.

Note – To view ALOM boot messages, you must connect an alphanumeric terminal to the serial management port *before* connecting the power cords to the server.

The sc> prompt indicates that you are interacting with the ALOM directly. It is the first prompt you see when you log in to the system through the serial management port or network management port, regardless of system power state.

Note – When you access the ALOM for the first time and you issue an administrative command, the system controller forces you to create a password (for the default username admin) for subsequent access. After this initial configuration, you are prompted to enter a user name and password every time you access the ALOM.

For more information about ALOM, see Chapter 2.

▼ To Obtain the ALOM Prompt From the Solaris Console

- **When connected to the Solaris console, type the escape sequence to take the console to the ALOM prompt.**

  By default the escape sequence is `#.` (hash-period).

  For instance, if the escape sequence is the default of `#.` you will type:

```
# #.
sc>
```

Note – Unlike the example, you will not see the `#.` being typed.

When you type the first character of the escape sequence, there is a one-second delay before the character appears on the screen. During this interval, you must type the second character of the escape sequence. If the escape sequence is completed within the one-second interval, the `sc>` prompt appears. Any characters typed after the second escape character are appended to the `sc>` prompt.

If the second escape character is incorrect, or is typed after the one-second interval has expired, then all characters are output at the original prompt.

▼ To Obtain the ALOM Prompt From the OpenBoot PROM

- **Type the sequence of escape characters.**

  By default, the escape sequence is `#.` (hash-period).

```
{2} ok #.
sc>
```

Note – Unlike the example, you will not see the `#.` being typed.

▼ To Connect to the Solaris Console From the ALOM Prompt

- **Use the `console` command from the ALOM prompt.**
  - If Solaris software is running, the system responds with the Solaris prompt:

```
sc>console
#
```

  - If the system was in the OpenBoot PROM, then the system responds with the OpenBoot PROM prompt:

```
sc>console
{2} ok
```

  - If the server is in Standby mode, the following message is generated:

```
sc>console
Solaris is not active
```

Note – The `console` command first attempts to connect to the Solaris console. If the Solaris console is not available, the `console` command then attempts to connect to the OpenBoot PROM. If the attempt is unsuccessful, the following message is displayed: `Solaris is not active`.

OpenBoot PROM `ok` Prompt

A server with the Solaris OS installed operates at different *run levels*. Most of the time, a server operates at run level 2 or run level 3, which are multiuser states with access to full system and network resources. Occasionally, you might operate the system at run level 1, which is a single-user administrative state. However, the lowest operational state is run level 0. At this state, it is safe to turn off power to the system.

When a server is at run level 0, the OpenBoot PROM `ok` prompt appears. This prompt indicates that the OpenBoot firmware is in control of the system.

For more information about the OpenBoot PROM, see Chapter 3.

▼ To Obtain the OpenBoot Prompt From the ALOM Prompt

- **Type the `break` command.**

```
sc> break
{2} ok
```

▼ To Obtain the OpenBoot Prompt When the Solaris OS Is Running

- **Type the `init 0` command at the Solaris prompt.**

```
# init 0
{1} ok
```

▼ To Terminate a Session When Connected to the System Controller Through the Serial Port

- If you are at the Solaris console or the OpenBoot PROM go to the ALOM prompt by typing the escape sequence, then terminate the ALOM prompt session by typing `logout` and pressing Return:

```
sc>logout
```

- If you are connected through a terminal server invoke the terminal server's command to disconnect the connection.
- If the connection was established using a `tip` command, then type the `tip` exit sequence `~.`(tilde and a period):

```
~.
```

▼ To Terminate a Session When Connected to the System Controller Through a Network Connection

1. **If you are at the Solaris prompt or the OpenBoot PROM, go to the ALOM prompt by typing the escape sequence.**
2. **Terminate the ALOM prompt session by using the `logout` command.**

   The remote session terminates automatically:

```
sc>logout
Connection closed by foreign host.
%
```

CHAPTER 2

Advanced Lights Out Manager

This chapter describes using Advanced Lights Out Manager (ALOM) for remote administration of your server. Topics include:

More information about ALOM is available in the *Advanced Lights Out Manager CMT v1.2 Guide*, 819-6672-10.

ALOM Overview

ALOM Features

ALOM is a system controller that comes preinstalled on your server, and is available as soon as you install and power on the system. Through a command-line interface, you can customize ALOM to your particular installation. Then you can monitor and control your server, either over the network or through a terminal server using the dedicated serial management port on the server.

What ALOM Monitors

TABLE 2-1 lists some of the components that ALOM can monitor on the server.

TABLE 2-1 What ALOM Monitors

Component Monitored	Information Provided
Disk drives	Whether each slot has a drive present, and whether the drive reports OK status
Fans	Fan speed and whether the fans report OK status
CPU temperatures	Whether a CPU is present, the temperature measured at the CPU, and any thermal warning or failure conditions
System enclosure temperature	System ambient temperature, as well as any enclosure thermal warning or failure conditions
Fuses	Whether fuses have been blown
Server front panel	System rotary switch position and status of LEDs
Voltages	Whether voltages are within operating range

Note – While redundant power sources are desirable, if only one DC connector is supplying power to the DC-powered version of the server, ALOM might occasionally report the following message:

```
SC Alert: env_log_event unsupported event
```

Using ALOM

The ALOM software is supplied ready-to-use and can support multiple users. However, only one user at a time can issue any commands that require write permissions. The other users can only issue read-only commands.

There are two ways to connect to ALOM:

- Use the `telnet` command to connect to ALOM through the Ethernet connection attached to the NET MGT port.
- Connect a serial device, such as an ASCII terminal or a port on a terminal server, to the SERIAL MGT port.

▼ To Set the Initial Password

When you first apply power to the server, ALOM automatically begins monitoring the system and displaying output to the system console using a preconfigured default account called `admin`, which has full (`cuar`) permissions. For security purposes, the administration password should be set.

1. **Physically connect to the ALOM serial management port and establish a connection.**

   Communication parameters are as follows:

   - 9600 Baud
   - 8 data bits
   - No parity
   - 1 stop bit
   - Full duplex
   - No handshaking

2. **Log in to the ALOM prompt. Type:**

```
#.
SC>
```

   That is:

   a. **Press and hold the Shift key and press the 3 key.**

   b. **Press the period key.**

   c. **Press the Return key.**

   The `sc>` prompt (ALOM prompt) is displayed.

3. **Type the** `password` **command.**

```
sc> password
```

4. **Type the password and re-type the password.**

   The password is created and is required for all future ALOM connections.

   If you do not log in before ALOM times out, ALOM reverts to the system console and displays the following message:

```
Enter #. to return to ALOM.
```

ALOM Shell Commands

The following tables list some of more common ALOM shell commands and briefly describes what these commands do.

Many ALOM shell commands can be executed from the Solaris command-line interface, using the `scadm` command. For example:

```
# scadm loghistory
```

Refer to the `scadm` man page for more information.

Configuration Commands

The ALOM configuration commands set or show the configuration of various aspects of the system.

TABLE 2-2 ALOM Configuration Commands

Command	Summary	Example
`password`	Changes the login password of the current user.	`sc> password`
`setdate` *mmddHHMMyyyy*	Sets the date and time, when the managed operating system is not running.	`sc> setdate 091321451999` `MON SEP 13 21:45:00 1999 UTC`
`setdefaults [-y] [-a]`	Resets all ALOM configuration parameters to their default values. The `-y` option enables you to skip the confirmation question. The `-a` option resets the user information to the factory default (one admin account only).	`sc> setdefaults -a`
`setsc` *parameter value*	Sets the specified ALOM *parameter* to the assigned *value*.	`sc> setsc netsc_ipaddr 1.2.3.4`
`setupsc`	Runs the interactive configuration script. This script configures the ALOM configuration variables.	`sc> setupsc`

TABLE 2-2 ALOM Configuration Commands *(Continued)*

Command	Summary	Example
`showdate`	Displays the ALOM set date. The Solaris OS and ALOM time are synchronized, but ALOM time is expressed in UTC (Coordinated Universal Time) rather than local time.	`sc>` **`showdate`** `MON SEP 13 21:45:00 1999 UTC`
`showplatform [-v]`	Displays information about the host system's hardware configuration, and whether the hardware is providing service. The `-v` option displays verbose information about the displayed component(s).	`sc>` **`showplatform`**
`showsc [-v]` *parameter*	Displays the current value of nonvolatile random access memory (NVRAM) configuration *parameters*. The `-v` option is needed for full version information.	`sc>` **`showsc sys_autorestart`** `xir`
`showusers [-g` *lines*`]`	Displays a list of users currently logged in to ALOM. The display for this command has a similar format to that of the UNIX command `who`. The `-g` option pauses the display after the number of lines you specify for *lines*.	`sc>` **`showusers -g 10`**
`useradd` *username*	Adds a user account to ALOM.	`sc>` **`useradd newuser`**
`userdel [-y]` *username*	Deletes a user account from ALOM. The `-y` option enables you to skip the confirmation question.	`sc>` **`userdel newuser`**
`userpassword` *username*	Sets or changes a user password.	`sc>` **`userpassword newuser`**
`userperm` *username* `[c][u][a][r]`	Sets the permission level for a user account.	`sc>` **`userperm newuser cr`**
`usershow [`*username*`]`	Displays a list of all user accounts, permission levels, and whether passwords are assigned.	`sc>` **`usershow newuser`**

FRU Commands

The ALOM FRU commands can show installed FRUs.

TABLE 2-3 ALOM FRU Commands

Command	Summary	Example
`removefru PS0|PS1`	Indicates if it is OK to perform a hot-swap of a power supply.	`sc>` **`removefru PS0`**
`showfru`	Displays information about the FRUs (field-replaceable units) in a host server.	`sc>` **`showfru`**

Log Commands

The ALOM log commands display the console output and ALOM event buffers.

TABLE 2-4 ALOM Log Commands

Command	Summary	Example
`consolehistory [-b` *lines*`|-e` *lines*`] [-g` *lines*`] [-v] [boot|run]`	Displays the host server console output buffers. The `-v` option displays the entire contents of the specified log.	`sc>` **`consolehistory boot -b 10`**
`showlogs [-b` *lines*`|-e` *lines*`] [-g` *lines*`] [-v]`	Displays the history of all events logged in the ALOM event buffer.	`sc>` **`showlogs -b 100`**

Status and Control Commands

The ALOM status and control commands enable you to perform typically manual tasks with the server, remotely.

TABLE 2-5 ALOM Status and Control Commands

Command	Summary	Example
`bootmode [skip_diag	diag	reset_nvram	normal	bootscript="`*string*`"]`	Controls the host server boot method through the OpenBoot PROM firmware.	`sc>` **`bootmode reset_nvram`** `sc>` **`reset`**
`break [-y] [-c]`	Drops the host server from the system into OpenBoot PROM or `kadb`.	`sc>` **`break`**
`clearasrdb`	Removes all entries from the `asr-db` blacklist.	`sc>` **`clearasrdb`**
`clearfault` *UUID*	Manually clears host-detected faults. *UUID* is the unique fault ID of the fault to be cleared.	`sc>` **`clearfault 1234`**
`console [-f]`	Connects to the host system console. The `-f` option forces the console write lock from one user to another.	`sc>` **`console`**
`disablecomponent` *asrkey*	Adds a component to the `asr-db` blacklist, where *asrkey* is the component to disable.	`sc>` **`disablecomponent MB/CMP0/CH3/R1/D1`**
`enablecomponent` *asrkey*	Removes a component from the asr-db blacklist, where *asrkey* is the component to enable.	`sc>` **`enablecomponent MB/CMP0/CH3/R1/D1`**
`flashupdate [-s` *IPaddr* `-f` *pathname*`] [-v]`	Updates the ALOM firmware. This command downloads `main` and `bootmon` firmware images to ALOM.	`sc>` **`flashupdate -s 1.2.3.4 -f /usr/platform/SUNW,Netra210/lib/images/alommainfw`**

TABLE 2-5 ALOM Status and Control Commands *(Continued)*

Command	Summary	Example
`powercycle [-f]`	Performs a `poweroff` followed by `poweron`. The `-f` option forces an immediate `poweroff`, otherwise the command attempts a graceful shutdown.	`sc> powercycle`
`poweroff [-y] [-f]`	Removes the main power from the host server. The `-y` option enables you to skip the confirmation question. The `-f` option forces an immediate shutdown.	`sc> poweroff`
`poweron [-c] [FRU]`	Applies the main power to the host server or a particular FRU.	`sc> poweron HDD1`
`reset [-y] [-x] [-c]`	Generates a hardware reset on the host server. The `-x` option generates an XIR (externally initiated reset). The `-y` option enables you to skip the confirmation question.	`sc> reset -x`
`setalarm critical|major|minor|user on|off`	Turns the alarm and associated LED on and off.	`sc> setalarm critical on`
`setkeyswitch [-y] normal | stby | diag | locked`	Sets the virtual keyswitch. The `-y` option enables you to skip the confirmation question when setting the keyswitch to `stby`.	`sc> setkeyswitch diag`
`setlocator on|off`	Turns the Locator LED on the server on or off. This function is available only on host servers that have Locator LEDs.	`sc> setlocator on`
`showcomponent`	Displays system components and their current state. The `showcomponent` command may not report all blacklisted DIMMS.	`sc> showcomponent`
`showfaults [-v]`	Displays current system faults. The `-v` option provides verbose output.	`sc> showfaults`
`showenvironment`	Displays the environmental status of the host server. This information includes system temperatures, power supply status, front panel LED status, hard drive status, fan status, voltage and current sensor status, and rotary switch position.	`sc> showenvironment`
`showkeyswitch`	Displays the status of the virtual keyswitch.	`sc> showkeyswitcht`
`showlocator`	Displays the current state of the Locator LED as either on or off. This function is available only on host servers that have Locator LEDs.	`sc> showlocator` `Locator LED is ON`
`shownetwork [-v]`	Displays the current network configuration information. The `-v` option shows additional information about your network, including information about your DHCP server.	`sc> shownetwork`

Other ALOM Commands

TABLE 2-6 lists other ALOM commands.

TABLE 2-6 Other ALOM Commands

Command	Summary	Example
`help`	Displays a list of all ALOM commands, or of a particular command, with their syntax and a brief description of how each command works.	`sc>` **`help poweron`**
`logout`	Logs out from an ALOM shell session.	`sc>` **`logout`**
`resetsc [-y]`	Reboots ALOM. The `-y` option enables you to skip the confirmation question.	`sc>` **`resetsc`**

Basic ALOM Tasks

Once you have logged in to ALOM as `admin` and specified the `admin` password, you can perform some common administrative tasks:

▼ To Reset ALOM

Resetting ALOM reboots the ALOM software. Reset ALOM after you have changed settings for ALOM or if ALOM stops responding for any reason.

- **At the `sc>` prompt, type `resetsc`.**

▼ To Switch Between the System Console and ALOM

- **To switch from the console to the ALOM** `sc>` **prompt, type** `#.` **(hash-period).**
- **To switch from the** `sc>` **prompt to the console, type** `console`**.**

▼ To Control the Locator LED

- **To turn the LED on and off, use the** `setlocator` **command.**
- **To check the state of the LED, use the** `showlocator` **command.**

  The LED can also be controlled at the Solaris superuser prompt using the `locator` command.

▼ To Reset the Host Server

1. **Type the** `poweroff` **command.**

   This message is displayed:

```
SC Alert: Host system has shut down.
```

2. **Type the** `poweron` **command.**

▼ To View Environmental Information About the Server

ALOM can display system temperatures, hard drive status, power supply and fan status, front panel LED status, rotary switch position, voltage and current sensors, alarm status, and so on.

- **To view environmental information, use the** `showenvironment` **command.**

▼ To Reconfigure ALOM to Use the Ethernet (NET MGT) Port

By default, ALOM uses the serial management port (SERIAL MGT) to communicate with a serial device. If desired, you can reconfigure ALOM to use the Ethernet network management (NET MGT) port, and then you can connect to ALOM through the `telnet` command.

Note – ALOM supports only 10-Mbit networks.

To configure the ALOM software to communicate using the NET MGT port, you must specify values for the network interface variables. The `setupsc` script helps you do this.

1. **Run the `setupsc` script. Type:**

```
sc> setupsc
```

The setup script starts. Answer the questions in the script. The script asks:

```
Do you wish to configure the enabled interfaces [y]?
```

2. **Type `y`.**

The script asks:

```
Should the SC network interface be enabled?
```

3. **Type `true` or press Return to enable the network interface.**

This sets a value for the `if_network` variable.

4. **Provide values for the following variables in the script:**
   - `if_modem` (specify `false`)
   - `netsc_dhcp` (`true` or `false`)
   - `netsc_ipaddr` (IP address)
   - `netsc_ipnetmask` (netmask)
   - `netsc_ipgateway` (IP address)
   - `netsc_tpelinktest` (`true` or `false`)

5. **When you have finished setting up the network interface variables, type Ctrl-Z to save your changes and exit the `setupsc` script.**

6. **Reset ALOM. Type:**

```
sc> resetsc
```

▼ To Add ALOM User Accounts

You can add a maximum of 15 unique user accounts to ALOM.

1. **Create an ALOM user account. Type:**

```
sc> useradd username
```

2. **Assign a password to this account. Type:**

```
sc> userpassword username
New password:
Re-enter new password:
```

3. **Assign permissions to this account. Type:**

```
sc> userperm username cuar
```

where *cuar* represents the `cuar` permissions.

4. **To verify accounts and their permissions, use the** `usershow` **command.**

▼ To Remove an ALOM User Account

- **To delete an ALOM user account, type:**

```
sc> userdel username
```

Note – You cannot delete the default `admin` account from ALOM.

▼ To Log In to ALOM

1. **Establish a connection with ALOM.**
2. **When the connection is established, type `#.` (hash-period) to escape from the system console.**
3. **Type in your ALOM login name and password.**

▼ To Change an ALOM Password

- **To change your password, use the `password` command.**
- **To change a user account password, use the `userpassword` *username* command.**

▼ To Set Up Email Alerts

Note – You can configure email alerts for up to eight users. You can configure each email address to receive its own severity level of alert.

---

1. **Ensure that ALOM is set up to use the Ethernet network management port (NET MGT), and that the network interface variables are configured.**

   See "To Reconfigure ALOM to Use the Ethernet (NET MGT) Port" on page 22.

2. **Configure email alerts and mail host. Type:**

```
sc> setsc if_emailalerts true
sc> setsc mgt_mailhost ipaddress1,...
```

3. **Configure each alert recipient. Type:**

```
sc> setsc mgt_mailalert emailaddress alertlevel
```

   where:

   - *emailaddress* is in the form of `emailusername@maildomain`
   - *alertlevel* is `1` for critical, `2` for major, and `3` for minor

4. **Repeat Step 3 for each alert recipient.**

   ALOM email alerts are displayed in the following format:

```
$HOSTID $EVENT $TIME $CUSTOMERINFO $HOSTNAME message
```

▼ To Back Up Your ALOM Configuration

You should periodically create a backup file on a remote system that records ALOM configuration settings.

- **As superuser, open a terminal window and type:**

```
# /usr/platform/SUNW,Netra210/sbin/scadm show > remote-filename
# /usr/platform/SUNW,Netra210/sbin/scadm usershow > remote-filename
```

Use a meaningful file name that includes the name of the server that ALOM controls. Later, you can refer to this file to restore the settings, if necessary.

▼ To Display Your ALOM Version

- **To display your ALOM version, type:**

```
sc> showsc version
Advanced Lights Out Manager v1.6
```

CHAPTER 3

OpenBoot PROM

This chapter describes the function, methods of obtaining, and configuration of the OpenBoot PROM. Topics include:

OpenBoot PROM Overview

The OpenBoot PROM is the low-level firmware that enables the server to boot into the Solaris Operating System. Once running Solaris, the OpenBoot PROM releases control of the server to the Solaris OS. Under certain conditions, the OpenBoot PROM will regain control of the server. A list of scenarios under which OpenBoot firmware control can occur follows:

- When you deliberately place the system under firmware control in order to execute firmware-based commands. This situation most often concerns you as an administrator, since there will be times when you need to obtain the `ok` prompt.
- By default, before the operating system is installed the system comes up under OpenBoot firmware control.
- When the `auto-boot?` OpenBoot variable is set to `false`, the system boots to the `ok` prompt.
- When the operating system is halted, the system transitions to run level 0 in an orderly way.
- When the operating system crashes, the system reverts to OpenBoot firmware control.

- During the boot process, when there is a serious hardware problem that prevents the operating system from running, the system reverts to OpenBoot firmware control.
- When a serious hardware problem develops while the system is running, the operating system transitions smoothly to run level 0.

Before You Obtain the `ok` Prompt

Note – Accessing the `ok` prompt suspends the Solaris OS. Before suspending the operating system, you should back up files, warn users of the impending shutdown, and bring the system down in an orderly manner.

Caution – When you access the `ok` prompt from a functioning server, you are suspending the Solaris OS and placing the system under firmware control. Any processes that were running under the operating system are also suspended, and *the state of such processes might not be recoverable*.

The commands you run from the `ok` prompt have the potential to affect the state of the system. This means that it is not always possible to resume execution of the operating system from the point at which it was suspended. Although the `go` command will resume execution in most circumstances, in general, each time you obtain the `ok` prompt, you should expect to have to reboot the system to get back to the operating system.

Obtaining the `ok` Prompt

There are several ways to obtain the `ok` prompt. In order of desirability, these are:

- Graceful shutdown
- ALOM `break` and `console` command
- Stop-A keys or Break key
- Manual system reset

Note – After forcing the system into OpenBoot firmware control, be aware that issuing certain OpenBoot commands (`probe-scsi`, `probe-scsi-all`, or `probe-ide`) might hang the system.

Graceful Shutdown

The preferred method of obtaining the `ok` prompt is to shut down the operating system by issuing an appropriate command (for example, the `shutdown`, `init`, or `uadmin` command) as described in Solaris system administration documentation. You can also use the system Power button to initiate a graceful system shutdown.

Gracefully shutting down the system prevents data loss, enables you to warn users beforehand, and causes minimal disruption. You can usually perform a graceful shutdown, provided the Solaris OS is running and the hardware has not experienced serious failure.

You can also perform a graceful system shutdown from the ALOM command prompt.

ALOM `break` or `console` Command

Typing `break` from the `sc>` prompt forces the server to change to OpenBoot firmware control. If the operating system is already halted, you can use the `console` command instead of `break` to reach the `ok` prompt.

Stop-A Keys or Break Key

When it is impossible or impractical to shut down the system gracefully, you can obtain the `ok` prompt by typing the Stop-A key sequence from a Sun keyboard. If you have an alphanumeric terminal attached to the server, press the Break key.

Note – These methods of reaching the `ok` prompt will only work if the system console has been redirected to the appropriate port.

Manual System Reset

Caution – Forcing a manual system reset results in loss of system state data, and should be attempted only as a last resort. After a manual system reset, all state information is lost, which inhibits troubleshooting the cause of the problem until the problem recurs.

Use the ALOM `reset` command, or `poweron` and `poweroff` commands, to reset the server. Using these commands results in the loss of all system coherence and state information. A manual system reset could corrupt the server's file systems, although the `fsck` command usually restores them. Use this method only when nothing else works.

▼ To Obtain the `ok` Prompt

1. **Decide which method you need to use to reach the `ok` prompt.**
2. **Follow the appropriate instructions in** TABLE 3-1**.**

TABLE 3-1 Methods of Obtaining the `ok` Prompt

Method	What to Do
Graceful shutdown of the Solaris OS	From a shell or command tool window, issue an appropriate command (for example, the `shutdown` or `init 0` command) as described in Solaris system administration documentation.
Stop-A keys or Break key	• From a Sun keyboard connected directly to the server, press the Stop and A keys simultaneously. • From an alphanumeric terminal configured to access the system console, press the Break key.
ALOM `break` and `console` commands	From the `sc>` prompt, type the `break` command. Then issue the `console` command, provided the operating system software is not running and the server is already under OpenBoot firmware control.
Manual system reset	1. From the `sc>` prompt, type: `sc>` **`bootmode bootscript="setenv auto-boot? false"`** 2. Press Enter. 3. Then type: `sc>` **`reset`**

OpenBoot PROM Configuration Variables

▼ To Change an OpenBoot PROM Configuration Variable

- **Use the `setenv` command.**

  For example:

```
ok setenv diag-switch? true
```

  This example enables diagnostics.

  TABLE 3-2 describes the OpenBoot firmware configuration variables stored in non-volatile memory on the system. The OpenBoot configuration variables are printed here in the order in which they appear when you issue the `showenv` command.

TABLE 3-2 OpenBoot Configuration Variables Stored on the System Configuration Card

Variable	Possible Values	Default Value	Description
`local-mac-address?`	`true, false`	`true`	If `true`, network drivers use their own MAC address, not the server MAC address.
`fcode-debug?`	`true, false`	`false`	If `true`, includes name fields for plug-in device FCodes.
`scsi-initiator-id`	`0-15`	`7`	SCSI ID of the serial attached SCSI controller.
`oem-logo?`	`true, false`	`false`	If `true`, uses custom OEM logo; otherwise, uses Sun logo.
`oem-banner?`	`true, false`	`false`	If `true`, uses custom OEM banner.
`ansi-terminal?`	`true, false`	`true`	If `true`, enables ANSI terminal emulation.
`screen-#columns`	`0-n`	`80`	Sets number of columns on screen.
`screen-#rows`	`0-n`	`34`	Sets number of rows on screen.

TABLE 3-2 OpenBoot Configuration Variables Stored on the System Configuration Card *(Continued)*

Variable	Possible Values	Default Value	Description
ttya-rts-dtr-off	true, false	false	If true, operating system does not assert rts (request-to-send) and dtr (data-transfer-ready) on serial management port.
ttya-ignore-cd	true, false	true	If true, operating system ignores carrier-detect on serial management port.
ttya-mode	9600,8,n,1,-	9600,8,n,1,-	Serial management port (baud rate, bits, parity, stop, handshake). The serial management port only works at the default values.
output-device	virtual-console, screen	virtual-console	Power-on output device.
input-device	virtual-console, keyboard	virtual-console	Power-on input device.
auto-boot-on-error?	true, false	false	If true, boots automatically after system error.
load-base	0-n	16384	Address.
auto-boot?	true, false	true	If true, boots automatically after power on or reset.
boot-command	*variable-name*	boot	Action following a boot command.
boot-file	*variable-name*	none	File from which to boot if diag-switch? is false.
boot-device	*variable-name*	disk net	Devices from which to boot if diag-switch? is false.
use-nvramrc?	true, false	false	If true, executes commands in NVRAMRC during server startup.
nvramrc	*variable-name*	none	Command script to execute if use-nvramrc? is true.
security-mode	none, command, full	none	Firmware security level.
security-password	*variable-name*	none	Firmware security password if security-mode is not none (never displayed). *Do not set this directly.*
security-#badlogins	*variable-name*	none	Number of incorrect security password attempts.

TABLE 3-2 OpenBoot Configuration Variables Stored on the System Configuration Card *(Continued)*

Variable	Possible Values	Default Value	Description
diag-switch?	true, false	false	If true: 1. OpenBoot verbosity is set to maximum 2. After a boot request, boot diag-file from diag-device If false: 1. OpenBoot verbosity is set to minimum 2. After a boot request, boot boot-file from boot-device
error-reset-recovery	boot, sync, none	boot	Command to execute following a system reset generated by an error.
network-boot-arguments	[*protocol*,] [*key=value*,]	none	Arguments to be used by the PROM for network booting. Defaults to an empty string. network-boot-arguments can be used to specify the boot protocol (RARP/DHCP) to be used and a range of system knowledge to be used in the process. For further information, see the eeprom (1M) man page or your Solaris reference manual.

OpenBoot Emergency Procedures

The introduction of Universal Serial Bus (USB) keyboards with the newest Sun systems has made it necessary to change some of the OpenBoot emergency procedures. Specifically, the Stop-N, Stop-D, and Stop-F commands that were available on systems with non-USB keyboards are not supported on systems that use USB keyboards. If you are familiar with the earlier (non-USB) keyboard functionality, this section describes the analogous OpenBoot emergency procedures available in newer systems that use USB keyboards.

Stop-A Functionality

Stop-A (Abort) key sequence works the same as it does on systems with standard keyboards, except that it does not work during the first few seconds after the server is reset. In addition, you can issue the ALOM break command. For more information, see "Switching Between the Consoles" on page 6.

Stop-N Functionality

Stop-N functionality is not available. However, the Stop-N functionality can be closely emulated by completing the following steps, provided the system console is configured to be accessible using either the serial management port or the network management port.

▼ To Restore OpenBoot Configuration Defaults

1. **Log in to the ALOM.**

   See "Switching Between the Consoles" on page 6.

2. **Type the following command:**

```
sc> bootmode reset_nvram
sc> bootmode bootscript="setenv auto-boot? false"
sc>
```

Note – If you do not issue the `poweroff` and `poweron` commands, or the `reset` command within 10 minutes, the host server ignores the `bootmode` command.

Issue the `bootmode` command without arguments to display the current setting.

```
sc> bootmode
Bootmode: reset_nvram
Expires WED SEP 09 09:52:01 UTC 2005
bootscript="setenv auto-boot? false"
```

3. **To reset the system, type the following command:**

```
sc> reset
Are you sure you want to reset the system [y/n]?  y
sc>
```

4. **To view console output as the system boots with default OpenBoot configuration variables, switch to `console` mode.**

```
sc> console

ok
```

5. **Type `set-defaults` to discard any customized IDPROM values and to restore the default settings for all OpenBoot configuration variables.**

Stop-F Functionality

The Stop-F functionality is not available on systems with USB keyboards.

Stop-D Functionality

The Stop-D (Diags) key sequence is not supported on systems with USB keyboards. However, the Stop-D functionality can be closely emulated by setting the virtual keyswitch to `diag`, using the ALOM `setkeyswitch` command.

CHAPTER 4

Basic Administrative Tasks

This chapter describes status indicators and basic tasks that might be done as the course of system administration. Topics include:

Status Indicators

The system has LED indicators associated with the server itself and with various components. The server status indicators are located on the bezel and repeated on the back panel. The components with LED indicators to convey status are the dry contact alarm card, power supply units, Ethernet port, and hard drives.

The topics in this section include:

Interpreting Status LEDs

The behavior of LEDs on the server conform to the American National Standards Institute (ANSI) Status Indicator Standard (SIS). These standard LED behaviors are described in TABLE 4-1.

TABLE 4-1 LED Behavior and Meaning

LED Behavior	Meaning
Off	The condition represented by the color is not true.
Steady on	The condition represented by the color is true.
Standby blink	The system is functioning at a minimal level and ready to resume full function.
Slow blink	Transitory activity or new activity represented by the color is taking place.
Fast blink	Attention is required.
Feedback flash	Activity is taking place commensurate with the flash rate (such as disk drive activity).

The LEDs have assigned meanings, described in TABLE 4-2.

TABLE 4-2 LED Behaviors With Assigned Meanings

Color	Behavior	Definition	Description
White	Off	Steady state	
	Fast blink	4Hz repeating sequence, equal intervals on and off.	This indicator helps you to locate a particular enclosure, board, or subsystem (for example, the Locator LED).
Blue	Off	Steady state	
	Steady on	Steady state	If blue is on, a service action can be performed on the applicable component with no adverse consequences (for example, the OK-to-Remove LED).
Yellow/Amber	Off	Steady state	
	Slow blink	1Hz repeating sequence, equal intervals on and off.	This indicator signals new fault conditions. Service is required (for example, the Service Required LED).
	Steady on	Steady state	The amber indicator stays on until the service action is completed and the system returns to normal function.

TABLE 4-2 LED Behaviors With Assigned Meanings *(Continued)*

Color	Behavior	Definition	Description
Green	Off	Steady state	
	Standby blink	Repeating sequence consisting of a brief (0.1 sec.) on flash followed by a long off period (2.9 sec.)	The system is running at a minimum level and is ready to be quickly revived to full function (for example, the System Activity LED).
	Steady on	Steady state	Status normal; system or component functioning with no service actions required
	Slow blink		A transitory (temporary) event is taking place for which direct proportional feedback is not needed or not feasible.

Bezel Server Status Indicators

FIGURE 4-1 shows the location of the bezel indicators, and TABLE 4-3 provides information about the server status indicators.

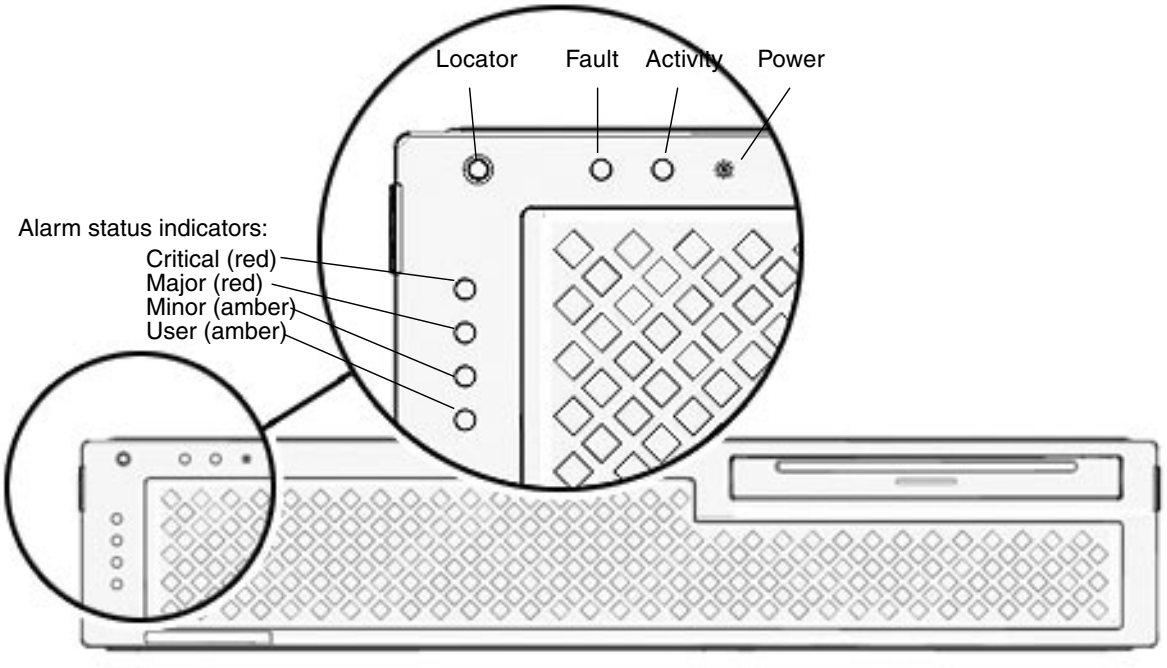

FIGURE 4-1 Location of the Bezel Server Status and Alarm Status Indicators

TABLE 4-3 Bezel Server Status Indicators

Indicator	LED Color	LED State	Component Status
Locator	White	On	Server is identified with the superuser `locator` or ALOM `setlocator` command.
		Off	Normal state
Fault	Amber	On	The server has detected a problem and requires the attention of service personnel.
		Off	The server has no detected faults.
Activity	Green	On	The server is powered up and running the Solaris Operating System.
		Off	Either power is not present or the Solaris software is not running.

You can check the status and turn the Locator LED on and off from either the superuser or ALOM prompt. TABLE 4-4 lists the commands.

TABLE 4-4 Locator LED Commands

Prompt	Status	Turn On	Turn Off
Superuser	`# /usr/sbin/locator`	`# /usr/sbin/locator -n`	`# /usr/sbin/locator -f`
ALOM	`sc> showlocator`	`sc> setlocator on`	`sc> setlocator off`

Alarm Status Indicators

The dry contact alarm card has four LED status indicators that are supported by ALOM. They are located vertically on the bezel (FIGURE 4-1). Information on the alarm indicators and dry contact alarm states is provided in TABLE 4-5. For more information on alarm indicators, see the *Advanced Lights Out Manager CMT v1.2 Guide,* 819-6672-10.

TABLE 4-5 Alarm Indicators and Dry Contact Alarm States

Indicator and Relay Labels	Indicator Color	Application or Server State	Condition or Action	Activity Indicator State	Alarm Indicator State	Relay NC[§] State	Relay NO[**] State	Comments
Critical (`Alarm0`)	Red	Server state (Power on or off, and Solaris OS functional or not functional)	No power input	Off	Off	Closed	Open	Default state
			System power off	Off	Off[‡]	Closed	Open	Input power connected
			System power turns on, Solaris OS not fully loaded	Off	Off[‡]	Closed	Open	Transient state
			Solaris OS successfully loaded	On	Off	Open	Closed	Normal operating state
			Watchdog timeout	Off	On	Closed	Open	Transient state, reboot Solaris OS
			Solaris OS shutdown initiated by user[*]	Off	Off[‡]	Closed	Open	Transient state
			Lost input power	Off	Off	Closed	Open	Default state
			System power shutdown by user	Off	Off[‡]	Closed	Open	Transient state
		Application state	User sets critical alarm to `on`[†]	--	On	Closed	Open	Critical fault detected
			User sets critical alarm to `off`[†]	--	Off	Open	Closed	Critical fault cleared
Major (`Alarm1`)	Red	Application state	User sets major alarm to `on`[†]	--	On	Open	Closed	Major fault detected
			User sets major alarm to `off`[†]	--	Off	Closed	Open	Major fault cleared

TABLE 4-5 Alarm Indicators and Dry Contact Alarm States *(Continued)*

Indicator and Relay Labels	Indicator Color	Application or Server State	Condition or Action	Activity Indicator State	Alarm Indicator State	Relay NC§ State	Relay NO** State	Comments
Minor (`Alarm2`)	Amber	Application state	User sets minor alarm to `on`†	--	On	Open	Closed	Minor fault detected
			User sets minor alarm to `off`†	--	Off	Closed	Open	Minor fault cleared
User (`Alarm3`)	Amber	Application state	User sets user alarm to `on`†	--	On	Open	Closed	User fault detected
			User sets user alarm to `off`†	--	Off	Closed	Open	User fault cleared

* The user can shut down the system using commands such as `init0` and `init6`. These commands do not remove power from the system.

† Based on a determination of the fault conditions, the user can turn the alarm on using the Solaris platform alarm API or ALOM CLI.

‡ The implementation of this alarm indicator state is subject to change.

§ NC state is the normally closed state. This state represents the default mode of the relay contacts in the normally closed state.

** NO state is the normally open state. This state represents the default mode of the relay contacts in the normally open state.

When the user sets an alarm, a message is displayed on the console. For example, when the critical alarm is set, the following message is displayed on the console:

```
SC Alert: CRITICAL ALARM is set
```

In certain instances when the critical alarm is set, the associated alarm indicator is not lit. This implementation is subject to change in future releases.

Selecting a Boot Device

The boot device is specified by the OpenBoot configuration variable `boot-device`. The default setting of this variable is `disk net`. With this setting, the firmware first attempts to boot from the system hard drive, and if that fails, from the on-board NET0 Gigabit Ethernet interface.

This procedure assumes that you are familiar with the OpenBoot firmware and that you know how to enter the OpenBoot environment. For more information, see "OpenBoot PROM" on page 27.

If you want to boot from a different device, perform the following procedure.

▼ To Select a Boot Device

1. **Obtain the `ok` prompt.**

   See "To Obtain the `ok` Prompt" on page 30.

2. **At the `ok` prompt, type:**

```
ok setenv boot-device device-specifier
```

   where *device-specifier* is one of the following:

   - `cdrom` – Specifies the optical media drive
   - `disk` – Specifies the system boot disk (internal disk 0 by default)
   - `disk0` – Specifies internal drive 0
   - `disk1` – Specifies internal drive 1
   - `disk2` – Specifies internal drive 2
   - `disk3` – Specifies internal drive 3
   - `net`, `net0`, `net1`, `net2`, `net3` – Specifies the network interfaces
   - *full path name* – Specifies the device or network interface by its full path name

---

Note – The Solaris OS modifies the `boot-device` variable to its full path name, not the alias name. If you choose a nondefault `boot-device` variable, the Solaris OS specifies the full device path of the boot device.

Note – You can specify the name of the program to be booted as well as the way the boot program operates. For more information, refer to the *OpenBoot 4.x Command Reference Manual* in the *OpenBoot Collection AnswerBook* for your specific Solaris OS release.

If you want to specify a network interface other than an on-board Ethernet interface as the default boot device, you can determine the full path name of each interface by typing:

```
ok show-devs
```

The `show-devs` command lists the system devices and displays the full path name of each PCI device.

Note – To boot off of a network interface, you must have a boot server available on the network.

Unconfiguring and Reconfiguring Devices

To support a degraded boot capability, the ALOM firmware provides the `disablecomponent` command, which enables you to unconfigure system devices manually. This command creates an entry in the ASR database, with the specified device flagged as disabled. Any device marked `disabled`, whether manually or by the system's firmware diagnostics, is removed from the system's machine description prior to the hand-off to other layers of system firmware, such as OpenBoot PROM.

▼ To Unconfigure a Device Manually

1. **Obtain the ALOM prompt.**

   See "Switching Between the Consoles" on page 6.

2. **At the `sc>` prompt, type:**

```
sc> disablecomponent asr-key
```

   where *asr-key* is one of the device identifiers from TABLE 4-6

Note – The device identifiers are not case-sensitive. You can type them as uppercase or lowercase characters.

TABLE 4-6 Device Identifiers and Devices

Device Identifiers	Devices
`MB/CMP`*cpu-number*`/P`*strand-number*	CPU strand (Number: 0-31)
`PCIE`*slot-number*	PCI-E slot (Number: 0-2)
`PCIX`*slot-number*	PCI-X (Number: 0-1):
`IOBD/PCIEa`	PCI-E leaf A (`/pci@780`)

TABLE 4-6 Device Identifiers and Devices *(Continued)*

Device Identifiers *(Continued)*	Devices *(Continued)*
`IOBD/PCIEb`	PCI-E leaf B (`/pci@7c0`)
`TTYA`	DB9 serial port
`MB/CMP0/CH`*channel-number*`/R`*rank-number*`/D`*dimm-number*	DIMMS

▼ To Reconfigure a Device Manually

1. **Obtain the ALOM prompt.**

   See "Switching Between the Consoles" on page 6.

2. **At the `sc>` prompt, type:**

```
sc> enablecomponent asr-key
```

   where *asr-key* is any device identifier from TABLE 4-6.

Note – The device identifiers are not case-sensitive. You can type them as uppercase or lowercase characters.

You can use the ALOM `enablecomponent` command to reconfigure any device that you previously unconfigured with the `disablecomponent` command.

Displaying System Fault Information

ALOM software lets you display current valid system faults. The `showfaults` command displays the fault ID, the faulted FRU device, and the fault message to standard output. The `showfaults` command also displays POST results.

▼ To Display System Fault Information

1. **Obtain the ALOM prompt.**

   See "Switching Between the Consoles" on page 6.

2. **At the** `sc>` **prompt type:**

```
sc> showfaults -v
```

For example:

```
sc> showfaults
   ID FRU       Fault
    0 FT0.FM2   SYS_FAN at FT0.FM2 has FAILED.
```

Adding the –v option displays the time:

```
sc> showfaults -v
   ID Time               FRU         Fault
    0 MAY 20 10:47:32   FT0.FM2     SYS_FAN at FT0.FM2 has FAILED.
```

Multipathing Software

Multipathing software enables you to define and control redundant physical paths to I/O devices, such as storage devices and network interfaces. If the active path to a device becomes unavailable, the software can automatically switch to an alternate path to maintain availability. This capability is known as *automatic failover*. To take advantage of multipathing capabilities, you must configure the server with redundant hardware, such as redundant network interfaces or two host bus adapters connected to the same dual-ported storage array.

For the server, three different types of multipathing software are available:

- Solaris IP Network Multipathing software provides multipathing and load-balancing capabilities for IP network interfaces.
- VERITAS Volume Manager (VxVM) software includes a feature called Dynamic Multipathing (DMP), which provides disk multipathing as well as disk load balancing to optimize I/O throughput.
- Sun StorEdge™ Traffic Manager is an architecture fully integrated within the Solaris OS (beginning with the Solaris 8 release) that enables I/O devices to be accessed through multiple host controller interfaces from a single instance of the I/O device.

For instructions on how to configure and administer Solaris IP Network Multipathing, consult the *IP Network Multipathing Administration Guide* provided with your specific Solaris release.

For information about VxVM and its DMP feature, refer to the documentation provided with the VERITAS Volume Manager software.

For information about Sun StorEdge Traffic Manager, refer to your Solaris OS documentation.

Storing FRU Information

▼ To Store Information in Available FRU PROMs

1. **Obtain the ALOM prompt.**

   See "Switching Between the Consoles" on page 6.

2. **At the `sc>` prompt type:**

```
setfru -c data
```

Automatic System Recovery

Automatic System Recovery (ASR) consists of self-test features, and an auto-configuring capability to detect failed hardware components and unconfigure them. By enabling this, the server can resume operating after certain nonfatal hardware faults or failures have occurred.

If a component is monitored by ASR and the server can operate without the components, the server automatically reboots if that component should develop a fault or fail. This functionality prevents a faulty hardware component from keeping the entire system down or causing the system to fail repeatedly.

If a fault is detected during the power-on sequence, the faulty component is disabled. If the system remains capable of functioning, the boot sequence continues.

To support this degraded boot capability, the OpenBoot firmware uses the 1275 Client Interface (by means of the device tree) to mark a device as either *failed* or *disabled*, by creating an appropriate status property in the device tree node. The Solaris Operating System does not activate a driver for any subsystem marked as failed or disabled.

As long as a failed component is electrically dormant (not causing random bus errors or signal noise, for example), the system reboots automatically and resumes operation while a service call is made.

Once a *failed* or *disabled* device is replaced with a new one, the OpenBoot firmware automatically modifies the status of the device upon reboot.

Note – ASR is not enabled until you activate it. See "To Enable ASR" on page 50.

Autoboot Options

The `auto-boot?` setting controls whether or not the firmware automatically boots the operating system after each reset. The default setting is `true`.

The `auto-boot-on-error?` setting controls whether the system attempts a degraded boot when a subsystem failure is detected. The default setting for `auto-boot-on-error?` is `false`. Both the `auto-boot?` and `auto-boot-on-error?` settings must be set to `true` to enable an automatic degraded boot.

▼ To Enable Automatic Degraded Boot

1. **Obtain the `ok` prompt.**

   See "To Obtain the `ok` Prompt" on page 30.

2. **Type:**

```
ok setenv auto-boot? true
ok setenv auto-boot-on-error? true
```

Note – The system does not attempt a degraded boot in response to any fatal non-recoverable error, even if degraded booting is enabled. For examples of fatal non-recoverable errors, see "Error Handling Summary" on page 49.

Error Handling Summary

Error handling during the power-on sequence falls into one of the following three cases:

- If no errors are detected by POST or OpenBoot diagnostics, the system attempts to boot if `auto-boot?` is `true`.
- If only nonfatal errors are detected by POST or OpenBoot diagnostics, the system attempts to boot if `auto-boot?` is `true` and `auto-boot-on-error?` is `true`. Nonfatal errors include the following:
  - SAS subsystem failure. In this case, a working alternate path to the boot disk is required. For more information, see "Multipathing Software" on page 47.
  - Ethernet interface failure.
  - USB interface failure.
  - Serial interface failure.
  - PCI card failure.
  - Memory failure. If a DIMM fails, the firmware will unconfigure the entire logical bank associated with the failed module. Another nonfailing logical bank must be present in the system for the system to attempt a degraded boot.

Note – If POST or OpenBoot diagnostics detects a nonfatal error associated with the normal boot device, the OpenBoot firmware automatically unconfigures the failed device and tries the next-in-line boot device, as specified by the `boot-device` configuration variable.

- If a fatal error is detected by POST or OpenBoot diagnostics, the system does not boot regardless of the settings of `auto-boot?` or `auto-boot-on-error?` Fatal nonrecoverable errors include the following:
  - All CPUs failed
  - All logical memory banks failed
  - Flash RAM cyclical redundancy check (CRC) failure
  - Critical field-replaceable unit (FRU) PROM configuration data failure
  - Critical application-specific integrated circuit (ASIC) failure

▼ To Enable ASR

1. **Obtain the `ok` prompt.**

   See "To Obtain the `ok` Prompt" on page 30.

2. **Configure the system for ASR. Type:**

```
ok setenv diag-switch? true
ok setenv auto-boot? true
ok setenv auto-boot-on-error? true
```

3. **Enable ASR. Type:**

```
ok reset-all
```

The system permanently stores the parameter changes and boots automatically.

▼ To Disable ASR

1. **Obtain the `ok` prompt.**

See "To Obtain the ok Prompt" on page 30.

2. **Unconfigure diagnostic modes. Type:**

```
ok setenv diag-switch? false
```

3. **Disable ASR. Type:**

```
ok reset-all
```

The system permanently stores the parameter changes and boots automatically.

Updating the Firmware

Updating or downgrading the firmware is accomplished using the `flashupdate` command from the ALOM prompt. The `flashupdate` command updates the flash PROMs in the system controller and the server motherboard. The `flashupdate` command requires the network management port to be connected to a suitable network. The network management port must be configured so that it can recognize an external FTP server that contains the new firmware images to be downloaded.

To use the `flashupdate` command, you need to know the following:

- IP address of the FTP server from which you want to download the firmware image
- Path at which the image is stored
- Username and password to enter at the prompts

If you do not have this information, ask your network administrator.

The syntax for the `flashupdate` command is:

`flashupdate [-s` *IPaddr* `-f` *pathname*`] [-v]`

where:

- `-s` *IPaddr* is the IP address of an FTP server having the firmware image
- `-f` *pathname* is the full directory path to the firmware image file
- `-v` enables verbose output of the download and update progression

---

Note – `flashupdate` cannot retrieve flash images from a secure (user ID and password) protected HTTP URL. A message of the form `flashupdate: failed, URL does not contain required file:` *file* is returned, although the file might exist.

Caution – Do not interrupt the `flashupdate` operation. If the `flashupdate` command is terminated abnormally, the system controller goes into single-user mode and is only accessible from the serial port.

▼ To Update the Server Firmware

1. **Power on the server.**
2. **Obtain the ALOM prompt.**

   See "Switching Between the Consoles" on page 6.
3. **Upgrade the firmware:**

```
sc> flashupdate -s IPaddr -f pathname
```

For example, (replace `123.45.67.89` with a valid IP address):

```
sc> flashupdate -s 123.45.67.89 -f
/net/server/sysfw/System_Firmware-6_0_0-Netra_T2000.bin

SC Alert: System poweron is disabled.
```

4. **When prompted, type your username and password.**

   For example:

```
Username: username
Password: password
```

   The username and password are based on your UNIX or LDAP user name and password, and not your ALOM username and password.

   After you type your username and password, the download process continues and a series of periods appear across your screen.

   For example:

```
.............................................................
.............................................................
.................................
```

   When the download process is finished, ALOM displays the message:

```
Update complete. Reset device to use new software.

SC Alert: SC firmware was reloaded
```

5. **Type the `resetsc` command to reset ALOM:**

```
sc> resetsc
Are you sure you want to reset the SC [y/n]? y
User Requested SC Shutdown
```

Note – To bypass the confirmation prompt, use the `-y` flag with the `resetsc` command. If `resetsc` is issued from a Telnet session, upon reset the Telnet session will be terminated. The output from the reset will be displayed through the serial management port of the system controller.

The system controller resets, runs diagnostics, and returns to the login prompt.

CHAPTER 5

Securing the Server

This chapter provides important information about securing the system, explains security recommendations, discusses domain minimization, and provides references to Solaris Operating System security.

This chapter includes the following topics:

Security Guidelines

The following are security practices to consider:

- Ensure that all passwords comply with security guidelines.
- Change your passwords on a regular basis.
- Scrutinize log files on a regular basis for any irregularities.

The practice of configuring a system to limit unauthorized access is called *hardening*. There are several configuration steps that can contribute to hardening your system. These steps are guidelines for system configuration:

- Implement security modifications immediately after updating the Sun Fire™ Real-Time Operating System (RTOS) and SC application firmware, and before configuring or installing any Sun Fire domains.
- In general, restrict access to the SC operating system, RTOS.
- Limit physical access to serial ports.
- Expect to reboot, depending upon the configuration changes.

Defining the Console Password

The only restrictions on SC console passwords are the character set supported by ASCII and the terminal emulator in use. The SC uses the MD5 algorithm to generate a hash of the password entered. Correspondingly, all characters entered are significant.

A minimum password length of 16 characters promotes the use of pass-phrases instead of passwords. Passwords should be composed of a mixture of lowercase, uppercase, numeric, and punctuation characters. For information on how to set the console password, see the *Netra T2000 Server Installation Guide*, 819-5838.

Using the SNMP Protocol Default Configuration

Simple Network Management Protocol (SNMP) is commonly used to monitor and manage networked devices and servers. By default, SNMP is disabled.

Note – The use of Sun Management Center software requires SNMP. However, since the SC does not support a secure version of the SNMP protocol, do not enable SNMP unless you must use Sun Management Center software.

Rebooting the System Controller to Implement Settings

▼ To Reboot the System Controller

The SC needs to be rebooted if a console message similar to the following is displayed:

```
Rebooting the SC is required for changes in network settings to
take effect.
```

1. **Type** `resetsc -y` **to reboot the SC.**

   The SC can be rebooted while the Solaris domain is running.

2. **Use the `shownetwork` command to validate that all the network modifications were implemented.**

   For information about using the Sun Security Toolkit to create secure configurations for servers running the Solaris Operating System, see the following web site:

   `http://www.sun.com/software/security/jass`

Selecting a Remote Connection Type

The SSH and Telnet services on the system controller are disabled by default.

Enabling SSH

If the system controller is on a general-purpose network, you can ensure secure remote access to the system controller by using SSH rather than Telnet. SSH encrypts data flowing between host and client. SSH provides authentication mechanisms that identify both hosts and users, enabling secure connections between known systems. Telnet is fundamentally insecure because the Telnet protocol transmits information, including passwords, unencrypted.

Note – SSH does not help with FTP, HTTP, SYSLOG, or SNMPv1 protocols. These protocols are unsecure and should be used cautiously on general purpose networks.

The system controller provides limited SSH functionality, supporting only SSH version 2 (SSHv2) client requests. TABLE 5-1 identifies the various SSH server attributes and describes how the attributes are handled in this subset. These attribute settings are not configurable.

TABLE 5-1 SSH Server Attributes

Attribute	Example Values	Comment
`Protocol`	`2`	Supports only SSH v2
`Port`	`22`	Listening port
`ListenAddress`	`0.0.0.0`	Supports multiple IP addresses
`AllowTcpForwarding`	`no`	Port forwarding not supported

TABLE 5-1 SSH Server Attributes *(Continued)*

Attribute	Example Values	Comment
RSAAuthentication	no	Public key authentication disabled
PubkeyAuthentication	no	Public key authentication disabled
PermitEmptyPasswords	yes	Password authentication controlled by the SC
MACs	hmac-sha1,hmac-md5	Same SSH server implementation as the Solaris 9 Operating System
Ciphers	aes128-cbc,blowfish-cbc,3des-cbc	Same SSH server implementation as the Solaris 9 Operating System

▼ To Enable SSH

- **To enable SSH, type:**

```
sc> setupsc
```

You are prompted to enter the network configuration and connection parameters.

For example:

```
sc> setupsc

Network Configuration
---------------------
Is the system controller on a network? [yes]:
Use DHCP or static network settings? [static]:
Hostname [hostname]:
IP Address [xxx.xxx.xxx.xxx]:
Netmask [xxx.xxx.xxx.x]:
Gateway [xxx.xxx.xxx.xxx]:
DNS Domain [xxxx.xxx.xxx]:
Primary DNS Server [xxx.xxx.xxx.xx]:
Secondary DNS Server [xxx.xxx.xx.x]:
Connection type (ssh, telnet, none) [ssh]:

Rebooting the SC is required for changes in the above network
settings to take effect.
lom>
```

Features Not Supported by SSH

The SSH server on the server does not support the following features:

- Remote command-line execution
- `scp` command (secure copy program)
- `sftp` command (secure file transfer program)
- Port forwarding
- Key-based user authentication
- SSH v1 clients

If you try to use any of the above features, an error message is generated. For example, if you type the following command:

```
# ssh SCHOST showboards
```

The following messages are generated:

- On the SSH client:

```
Connection to SCHOST closed by remote host.
```

- On the SC console:

```
[0x89d1e0] sshdSessionServerCreate: no server registered
           for showboards
[0x89d1e0] sshd: Failed to create sshdSession
```

Changing SSH Host Keys

It is good security practice to obtain new host keys periodically. If you suspect that the host key might be compromised, you can use the `ssh-keygen` command to regenerate system host keys.

Host keys, once generated, can only be replaced and not deleted without resorting to the `setdefaults` command. For newly generated host keys to be activated, the SSH server must be restarted either by running the `restartssh` command or through a reboot. For further information on the `ssh-keygen` and `restartssh` commands (with examples), see the *Sun Fire Entry-Level Midrange System Controller Command Reference Manual*, 819-1268.

Note – You can also use the `ssh-keygen` command to display the host key fingerprint on the system controller.

Additional Security Considerations

Special Key Sequences for RTOS Shell Access

Special key sequences can be issued to the SC, over its serial connection, while it is booting. These key sequences have special capabilities if entered at the serial port within the first 30 seconds after an SC reboot.

The special capabilities of these key sequences are automatically disabled 30 seconds after the Sun copyright message is displayed. Once the capability is disabled, the key sequences operate as normal control keys.

Because of the risk that the security of the SC could be compromised by unauthorized access to the RTOS shell, you should control access to the serial ports of the SC.

Domain Minimization

One way to contribute to the security of a server is to tailor the installation of software to an essential minimum. By limiting the number of software components installed on each domain (called *domain minimization*), you can reduce the risks of security holes that can be exploited by potential intruders.

For a detailed discussion of minimization, with examples, see *Minimizing Domains for Sun Fire V1280, 6800, 12K, and 15K Systems* (two-part article) available online at:

`http://www.sun.com/security/blueprints`

Solaris Operating System Security

For information on securing the Solaris Operating System, see the following books and articles:

- *Solaris Security Best Practices* – available online at:

  `http://www.sun.com/software/security/blueprints`

- *Solaris Security Toolkit* – available online at:

  `http://www.sun.com/software/security/jass`

CHAPTER 6

Managing Disk Volumes

This chapter describes redundant array of independent disks (RAID) concepts, and how to configure and manage RAID disk volumes using the server's on-board serial attached SCSI (SAS) disk controller.

This chapter discusses the following topics:

RAID Requirements

To configure and use RAID disk volumes on the server, you must install patch IDs 119850-12 and 122165-01. Patches are available for download from

`http://www.sunsolve.com`

Installation procedures for patches are included in text README files that accompany the patches.

Note – For the latest information on patches for the server, see the server product notes, available at: `http://www.sun.com/documentation`

Disk Volumes

From the perspective of the server's on-board disk controller, *disk volumes* are logical disk devices comprising one or more complete physical disks.

Once you create a volume, the operating system uses and maintains the volume as if it were a single disk. By providing this logical volume management layer, the operating system overcomes the restrictions imposed by physical disk devices.

The on-board disk controller of the server provides for the creation of as many as two hardware RAID volumes. The controller supports either two-disk RAID 1 (integrated mirror, or IM) volumes, or two-, three- or four-disk RAID 0 (integrated stripe, or IS) volumes.

Note – Due to the volume initialization that occurs on the disk controller when a new volume is created, properties of the volume such as geometry and size are unknown. RAID volumes created using the hardware controller must be configured and labeled using `format`(1M) prior to use with the Solaris Operating System. See "To Configure and Label a RAID Volume" on page 73, or the `format` (1M) man page for further details.

Volume migration (relocating all RAID volume disk members from one chassis to another) is not supported. If this operation must be performed, contact Sun Service.

RAID Technology

RAID technology allows for the construction of a logical volume, made up of several physical disks, to provide data redundancy, increased performance, or both. The server's on-board disk controller supports both RAID 0 and RAID 1 volumes.

This section describes the RAID configurations supported by the on-board disk controller:

- Integrated stripe, or IS volumes (RAID 0)
- Integrated mirror, or IM volumes (RAID 1)

Integrated Stripe Volumes (RAID 0)

Integrated Stripe volumes are configured by initializing the volume across two or more physical disks, and sharing the data written to the volume across each physical disk in turn, or *striping* the data across the disks.

Integrated stripe volumes provide for a logical unit (LUN) that is equal in capacity to the sum of all its member disks. For example, a three-disk IS volume configured on 72 GByte drives will have a 216 GByte capacity.

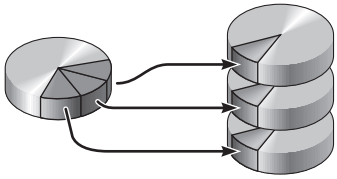

FIGURE 6-1 Graphical Representation of Disk Striping

Caution – There is no data redundancy in an IS volume configuration. Thus, if a single disk fails, the entire volume fails, and all data is lost. If an IS volume is manually deleted, all data on the volume is lost.

IS volumes are likely to provide better performance than IM volumes or single disks. Under certain workloads, particularly some write or mixed read-write workloads, I/O operations complete faster because each sequential block is written to each member disk in turn.

Integrated Mirror Volumes (RAID 1)

Disk mirroring (RAID 1) is a technique that uses data redundancy, where two complete copies of all data are stored on two separate disks to protect against loss of data due to disk failure. One logical volume is duplicated on two separate disks.

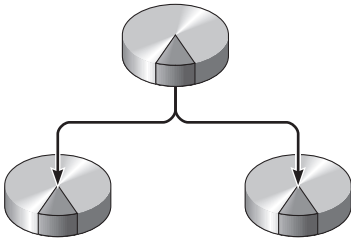

FIGURE 6-2 Graphical Representation of Disk Mirroring

Whenever the operating system needs to write to a mirrored volume, both disks are updated. The disks are maintained at all times with exactly the same information. When the operating system needs to read data, it reads from whichever disk is more readily accessible at the moment, which can result in enhanced performance for read operations.

Caution – Creating RAID volumes using the on-board disk controller destroys all data on the member disks. The disk controller's volume initialization procedure reserves a portion of each physical disk for metadata and other internal information used by the controller. Once the volume initialization is complete, you can configure the volume and label it using `format`(1M). You can then use the volume in the Solaris Operating System.

Hardware RAID Operations

On the server, the SAS controller supports mirroring and striping using the Solaris OS `raidctl` utility.

A hardware RAID volume created under the `raidctl` utility behaves slightly differently than one created using volume management software. Under a software volume, each device has its own entry in the virtual device tree, and read/write operations are performed to both virtual devices. Under hardware RAID volumes, only one device appears in the device tree. Member disk devices are invisible to the operating system, and are accessed only by the SAS controller.

Slot Numbers and Device Names for Non-RAID Disks

To perform a disk hot-swap procedure, you must know the physical or logical device name for the drive that you want to install or remove. If your system encounters a disk error, often you can find messages about failing or failed disks in the system console. This information is also logged in the `/var/adm/messages` files.

These error messages typically refer to a failed hard drive by its physical device name (such as `/devices/pci@1f,700000/scsi@2/sd@1,0`) or by its logical device name (such as `c0t1d0`). In addition, some applications might report a disk slot number (0 through 3).

You can use TABLE 6-1 to associate internal disk slot numbers with the logical and physical device names for each hard drive.

TABLE 6-1 Disk Slot Numbers, Logical Device Names, and Physical Device Names

Disk Slot Number	Logical Device Name*	Physical Device Name
Slot 0	`c0t0d0`	`/devices/pci@780/pci@0/pci@9/scsi@0/sd@0,0`
Slot 1	`c0t1d0`	`/devices/pci@780/pci@0/pci@9/scsi@0/sd@1,0`

* The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

▼ To Create a Mirrored Volume

1. **Verify which hard drive corresponds with which logical device name and physical device name.**

   See "Slot Numbers and Device Names for Non-RAID Disks" on page 67.

   To verify the current hardware RAID configuration, type:

```
# raidctl
No RAID volumes found.
```

The preceding example indicates that no RAID volume exists. In another case:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------
c0t0d0  IM      OK              c0t0d0          OK
                                c0t1d0          OK
```

In this example, a single IM volume has been enabled. It is fully synchronized and is online.

The server's on-board SAS controller can configure as many as two RAID volumes. Prior to volume creation, ensure that the member disks are available and that there are not two volumes already created.

The values provided in the `RAID Status` column are described as follows:

- `OK` – The RAID volume is online and fully synchronized
- `RESYNCING` – The data between the primary and secondary member disks in an IM are still synchronizing.
- `DEGRADED` – A member disk has failed or is otherwise offline.
- `FAILED` – The volume should be deleted and reinitialized. This can occur when any member disk in an IS volume is lost, or when both disks are lost in an IM volume.

The values provided in the `Disk Status` column are described as follows:

- `OK` – The drive is online and functioning properly
- `FAILED`, `MISSING`, or `OFFLINE` – The disk has hardware or configuration issues that need to be addressed.

For example, an IM with a secondary disk that has been removed from the chassis appears as:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------
c0t0d0  IM      DEGRADED        c0t0d0          OK
                                c0t1d0          MISSING
```

See the `raidctl`(1M) man page for additional details regarding volume and disk status.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

2. **Type the following command:**

```
# raidctl -c primary secondary
```

The creation of the RAID volume is interactive, by default. For example:

```
# raidctl -c c0t0d0 c0t1d0
Creating RAID volume c0t0d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume 'c0t0d0' created
#
```

As an alternative, you can use the `-f` option to force the creation if you are sure of the member disks, and sure that the data on both member disks can be lost. For example:

```
# raidctl -f -c c0t0d0 c0t1d0
Volume 'c0t0d0' created
#
```

When you create a RAID mirror, the secondary drive (in this case, `c0t1d0`) disappears from the Solaris device tree.

3. **(Optional) To check the status of a RAID mirror, type the following command:**

```
# raidctl
RAID    Volume  RAID          RAID          Disk
Volume  Type    Status        Disk          Status
-----------------------------------------------------------
c0t0d0  1M      RESYNCING     c0t0d0        OK
                              c0t1d0        OK
```

The preceding example indicates that the RAID mirror is still resynchronizing with the backup drive.

The following example shows that the RAID mirror is synchronized and online.

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------
c0t0d0  IM      OK              c0t0d0          OK
                                c0t1d0          OK
```

The disk controller synchronizes IM volumes one at a time. If you create a second IM volume before the first IM volume completes its synchronization, the first volume's RAID status will indicate `RESYNCING`, and the second volume's RAID status will indicate `OK`. Once the first volume has completed, its RAID status changes to `OK`, and the second volume automatically starts synchronizing, with a RAID status of `RESYNCING`.

Under RAID 1 (disk mirroring), all data is duplicated on both drives. If a disk fails, replace it with a working drive and restore the mirror. For instructions, see "To Perform a Mirrored Disk Hot-Swap Operation" on page 78.

For more information about the `raidctl` utility, see the `raidctl`(1M) man page.

▼ To Create a Mirrored Volume of the Default Boot Device

Due to the volume initialization that occurs on the disk controller when a new volume is created, the volume must be configured and labeled using the `format`(`1M`) utility prior to use with the Solaris Operating System (see "To Configure and Label a RAID Volume" on page 73). Because of this limitation, `raidctl`(`1M`) blocks the creation of a hardware RAID volume if any of the member disks currently have a file system mounted.

This section describes the procedure required to create a hardware RAID volume containing the default boot device. Since the boot device always has a mounted file system when booted, an alternate boot medium must be employed, and the volume created in that environment. The suggested alternate medium is a network installation image in single user mode. Refer to the *Solaris 10 Installation Guide* for information about configuring and using network-based installations.

1. **Determine which disk is the default boot device.**

   From the OpenBoot `ok` prompt, invoke the `printenv` command, and if necessary the `devalias` command, to identify the default boot device. For example:,

```
ok printenv boot-device
boot-device =         disk

ok devalias disk
disk                  /pci@780/pci@0/pci@9/scsi@0/disk@0,0
```

2. **Execute the `boot net -s` command.**

```
ok boot net -s
```

3. **Once the system has booted, use the `raidctl(1M)` utility to create a hardware mirrored volume, using the default boot device as the primary disk.**

   See "To Create a Mirrored Volume" on page 67. For example:,

```
# raidctl -c c0t0d0 c0t1d0
Creating RAID volume c0t0d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume  c0t0d0  created
#
```

   The volume can now be installed with the Solaris Operating System using any supported method. The hardware RAID volume `c0t0d0` appears as a disk to the Solaris installation program.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

▼ To Create a Striped Volume

1. **Verify which hard drive corresponds with which logical device name and physical device name.**

   See "Slot Numbers and Device Names for Non-RAID Disks" on page 67.

2. **(Optional) To verify the current RAID configuration, type:**

```
# raidctl
No RAID volumes found.
```

The preceding example indicates that no RAID volume exists.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

---

3. **Type the following command:**

```
# raidctl -c -r 0 disk1 disk2 ...
```

The creation of the RAID volume is interactive, by default. For example:

```
# raidctl -c -r 0 c0t1d0 c0t2d0 c0t3d0
Creating RAID volume c0t1d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume 'c0t1d0' created
#
```

When you create a RAID striped volume, the other member drives (in this case, `c0t2d0` and `c0t3d0`) disappear from the Solaris device tree.

As an alternative, you can use the `-f` option to force the creation if you are sure of the member disks, and sure that the data on all other member disks can be lost. For example:

```
# raidctl -f -c -r 0 c0t1d0 c0t2d0 c0t3d0
Volume 'c0t1d0' created
#
```

4. **(Optional) To check the status of a RAID striped volume, type the following command:**

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------------
c0t1d0  IS      OK              c0t1d0          OK
                                c0t2d0          OK
                                c0t3d0          OK
```

The example shows that the RAID striped volume is online and functioning.

Under RAID 0 (disk striping), there is no replication of data across drives. The data is written to the RAID volume across all member disks in a round-robin fashion. If any one disk is lost, all data on the volume is lost. For this reason, RAID 0 cannot be used to ensure data integrity or availability, but can be used to increase write performance in some scenarios.

For more information about the `raidctl` utility, see the `raidctl`(1M) man page.

▼ To Configure and Label a RAID Volume

After a creating a RAID volume using `raidctl`, use `format`(`1M`) to configure and label the volume before attempting to use it in the Solaris Operating System.

1. **Start the `format` utility.**

```
# format
```

The `format` utility might generate messages about corruption of the current label on the volume (which you are going to change). You can safely ignore these messages.

2. **Select the disk name that represents the RAID volume that you have configured.**
   In this example, `c0t2d0` is the logical name of the volume.

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
        0. c0t0d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@0,0
        1. c0t1d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@1,0
        2. c0t2d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@2,0
Specify disk (enter its number): 2
selecting c0t2d0
[disk formatted]
FORMAT MENU:
         disk       - select a disk
         type       - select (define) a disk type
         partition  - select (define) a partition table
         current    - describe the current disk
         format     - format and analyze the disk
         fdisk      - run the fdisk program
         repair     - repair a defective sector
         label      - write label to the disk
         analyze    - surface analysis
         defect     - defect list management
         backup     - search for backup labels
         verify     - read and display labels
         save       - save new disk/partition definitions
         inquiry    - show vendor, product and revision
         volname    - set 8-character volume name
         !<cmd>     - execute <cmd>, then return
         quit
```

3. **Issue the `type` command at the `format>` prompt, then select 0 (zero) to autoconfigure the volume.**

   For example:

```
format> type

AVAILABLE DRIVE TYPES:
         0. Auto configure
         1. DEFAULT
         2. SUN72G
         3. SUN72G
         4. other
Specify disk type (enter its number)[3]: 0
c0t2d0: configured with capacity of 68.23GB
<LSILOGIC-LogicalVolume-3000 cyl 69866 alt 2 hd 16 sec 128>
selecting c0t2d0
[disk formatted]
```

4. **Use the `partition` command to partition, or *slice*, the volume according to your desired configuration.**

   See the `format`(`1M`) man page for additional details.

5. **Write the new label to the disk using the `label` command.**

```
format> label
Ready to label disk, continue? yes
```

6. **Verify that the new label has been written by printing the disk list using the `disk` command.**

```
format> disk

AVAILABLE DISK SELECTIONS:
        0. c0t0d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@0,0
        1. c0t1d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@1,0
       2. c0t2d0 <LSILOGIC-LogicalVolume-3000 cyl 69866 alt 2 hd
16 sec 128>
           /pci@780/pci@0/pci@9/scsi@0/sd@2,0
Specify disk (enter its number)[2]:
```

Note – `c0t2d0` now has a type indicating it is an `LSILOGIC-LogicalVolume`.

7. **Exit the `format` utility.**

   The volume can now be used in the Solaris Operating System.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

▼ To Delete a RAID Volume

1. **Verify which hard drive corresponds with which logical device name and physical device name.**

   See "Slot Numbers and Device Names for Non-RAID Disks" on page 67.

2. **Determine the name of the RAID volume. Type the following command:**

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------
c0t0d0  IM      OK              c0t0d0          OK
                                c0t1d0          OK
```

In this example, the RAID volume is `c0t1d0`.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

3. **To delete the volume, type the following command:**

```
# raidctl -d mirrored-volume
```

For example:

```
# raidctl -d c0t0d0
RAID Volume ‘c0t0d0’ deleted
```

If the RAID volume is an IS volume, the deletion of the RAID volume is interactive, for example:

```
# raidctl -d c0t0d0
Deleting volume c0t0d0 will destroy all data it contains, proceed
(yes/no)? yes
Volume ’c0t0d0’ deleted.
#
```

The deletion of an IS volume results in the loss of all data that it contains. As an alternative, you can use the `-f` option to force the deletion if you are sure that you no longer need the IS volume, or the data it contains. For example:

```
# raidctl -f -d c0t0d0
Volume ’c0t0d0’ deleted.
#
```

4. **To confirm that you have deleted the RAID array, type the following command:**

```
# raidctl
```

For example:

```
# raidctl
No RAID volumes found
```

For more information, see the `raidctl`(1M) man page.

▼ To Perform a Mirrored Disk Hot-Swap Operation

1. **Verify which hard drive corresponds with which logical device name and physical device name.**

   See "Slot Numbers and Device Names for Non-RAID Disks" on page 67.

   If the Disk Status is FAILED, then the drive can be removed and a new drive inserted. Upon insertion, the new disk status should be OK and the volume should be RESYNCING.

2. **To confirm a failed disk, type the following command:**

```
# raidctl
```

   For example:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------------
c0t1d0  IM      DEGRADED        c0t1d0          OK
                                c0t2d0          FAILED
```

   This example indicates that the disk mirror has degraded due to a failure in disk `c0t2d0`.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

---

3. **Remove the hard drive, as described in the server service manual.**

   There is no need to issue a software command to bring the drive offline when the drive has failed.

4. **Install a new hard drive, as described in the server service manual.**

   The RAID utility automatically restores the data to the disk.

5. **To check the status of a RAID rebuild, type the following command:**

```
# raidctl
```

For example:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------------
c0t1d0  IM      RESYNCING       c0t1d0          OK
                                c0t2d0          OK
```

This example indicates that RAID volume `c0t1d0` is resynchronizing.

If you issue the command again once synchronization has completed, it indicates that the RAID mirror is finished resynchronizing and is back online:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------------
c0t1d0  IM      OK              c0t1d0          OK
                                c0t2d0          OK
```

For more information, see the `raidctl`(1M) man page.

▼ To Perform a Nonmirrored Disk Hot-Swap Operation

1. **Verify which hard drive corresponds with which logical device name and physical device name.**

   See "Slot Numbers and Device Names for Non-RAID Disks" on page 67.

   Ensure that no applications or processes are accessing the hard drive.

2. **View the status of the SCSI devices.**

   To view the status of the SCSI devices, type the following command:

```
# cfgadm -al
```

   For example:

```
# cfgadm -al
Ap_Id           Type         Receptacle   Occupant      Condition
c0              scsi-bus     connected    configured    unknown
c0::dsk/c0t0d0  disk         connected    configured    unknown
c0::dsk/c0t1d0  disk         connected    configured    unknown
c0::dsk/c0t2d0  disk         connected    configured    unknown
c0::dsk/c0t3d0  disk         connected    configured    unknown
c1              scsi-bus     connected    configured    unknown
c1::dsk/c1t0d0  CD-ROM       connected    configured    unknown
usb0/1          unknown      empty        unconfigured  ok
usb0/2          unknown      empty        unconfigured  ok
usb1/1.1        unknown      empty        unconfigured  ok
usb1/1.2        unknown      empty        unconfigured  ok
usb1/1.3        unknown      empty        unconfigured  ok
usb1/1.4        unknown      empty        unconfigured  ok
usb1/2          unknown      empty        unconfigured  ok
#
```

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

The `-al` options return the status of all SCSI devices, including buses and USB devices. (In this example, no USB devices are connected to the system.)

Note that while you can use the Solaris OS `cfgadm install_device` and `cfgadm remove_device` commands to perform a hard drive hot-swap procedure, these commands issue the following warning message when you invoke these commands on a bus containing the system disk:

```
# cfgadm -x remove_device c0::dsk/c0t1d0
Removing SCSI device: /devices/pci@1f,4000/scsi@3/sd@1,0
This operation will suspend activity on SCSI bus: c0
Continue (yes/no)? y
dev = /devices/pci@780/pci@0/pci@9/scsi@0/sd@1,0
cfgadm: Hardware specific failure: failed to suspend:
     Resource              Information
------------------  -------------------------
/dev/dsk/c0t0d0s0   mounted filesystem "/"
/dev/dsk/c0t0d0s6   mounted filesystem "/usr"
```

This warning is issued because these commands attempt to quiesce the (SAS) SCSI bus, but the server firmware prevents it. This warning message can be safely ignored in the server, but the following step avoids this warning message altogether.

3. **Remove the hard drive from the device tree.**

   To remove the hard drive from the device tree, type the following command:

```
# cfgadm -c unconfigure Ap-Id
```

   For example:

```
# cfgadm -c unconfigure c0::dsk/c0t3d0
```

   This example removes `c0t3d0` from the device tree. The blue OK-to-Remove LED lights.

4. **Verify that the device has been removed from the device tree.**

   To verify that the device has been removed from the device tree, type the following command:

```
# cfgadm -al
Ap_Id           Type         Receptacle   Occupant      Condition
c0              scsi-bus     connected    configured    unknown
c0::dsk/c0t0d0  disk         connected    configured    unknown
c0::dsk/c0t1d0  disk         connected    configured    unknown
c0::dsk/c0t2d0  disk         connected    configured    unknown
c0::dsk/c0t3d0  unavailable  connected    configured    unknown
c1              scsi-bus     connected    unconfigured  unknown
c1::dsk/c1t0d0  CD-ROM       connected    configured    unknown
usb0/1          unknown      empty        unconfigured  ok
usb0/2          unknown      empty        unconfigured  ok
usb1/1.1        unknown      empty        unconfigured  ok
usb1/1.2        unknown      empty        unconfigured  ok
usb1/1.3        unknown      empty        unconfigured  ok
usb1/1.4        unknown      empty        unconfigured  ok
usb1/2          unknown      empty        unconfigured  ok
#
```

   Note that `c0t3d0` is now `unavailable` and `unconfigured`. The corresponding hard drive OK-to-Remove LED is lit.

5. **Remove the hard drive, as described in the server service manual.**

   The blue OK-to-Remove LED goes out when you remove the hard drive.

6. **Install a new hard drive, as described in the server service manual.**

7. **Configure the new hard drive.**

   To configure the new hard drive, type the following command:

```
# cfgadm -c configure Ap-Id
```

   For example:

```
# cfgadm -c configure c1::dsk/c0t3d0
```

   The green Activity LED flashes as the new disk at `c1t3d0` is added to the device tree.

8. **Verify that the new hard drive is in the device tree.**

   To verify that the new hard drive is in the device tree, type the following command:

```
# cfgadm -al
Ap_Id           Type         Receptacle   Occupant      Condition
c0              scsi-bus     connected    configured    unknown
c0::dsk/c0t0d0  disk         connected    configured    unknown
c0::dsk/c0t1d0  disk         connected    configured    unknown
c0::dsk/c0t2d0  disk         connected    configured    unknown
c0::dsk/c0t3d0  disk         connected    configured    unknown
c1              scsi-bus     connected    configured    unknown
c1::dsk/c1t0d0  CD-ROM       connected    configured    unknown
usb0/1          unknown      empty        unconfigured  ok
usb0/2          unknown      empty        unconfigured  ok
usb1/1.1        unknown      empty        unconfigured  ok
usb1/1.2        unknown      empty        unconfigured  ok
usb1/1.3        unknown      empty        unconfigured  ok
usb1/1.4        unknown      empty        unconfigured  ok
usb1/2          unknown      empty        unconfigured  ok
#
```

   `c0t3d0` is now listed as configured.

APPENDIX **A**

Watchdog Timer Application Mode

This appendix gives information on the watchdog timer application mode on the server. It provides the following sections to help you understand how to configure and use the watchdog timer and to program Alarm3:

Note – Once the application watchdog timer is in use, it is necessary to reboot the Solaris operating system in order to return to the default (non-programmable) watchdog timer and default LED behavior (no Alarm3).

Understanding the Watchdog Timer Application Mode

The watchdog mechanism detects a system hang, or an application hang or crash, should they occur. The watchdog is a timer that is continually reset by a user application as long as the operating system and user application are running.

When the application is rearming the application watchdog, an expiration can be caused by:

- Crash of the rearming application
- Hang or crash of the rearming thread in the application
- System hang

When the system watchdog is running, a system hang, or more specifically, the hang of the clock interrupt handler causes an expiration.

The system watchdog mode is the default. If the application watchdog is not initialized, then the system watchdog mode is used.

The application mode enables you to:

- Configure the watchdog timer – Your applications running on the host can configure and use the watchdog timer, enabling you to detect fatal problems from applications and to recover automatically.
- Program Alarm3 – This capability enables you to generate this alarm in case of critical problems in your applications.

The `setupsc` command, an existing command for ALOM, can be used to configure the recovery for the system watchdog *only*:

```
sc> setupsc
```

The system controller configuration should be as follows:

```
SC POST diag Level [off]:
Host Watchdog [enabled]:
Rocker Switch [enabled]:
Secure Mode [off]:

PROC RTUs installed: 0
PROC Headroom quantity (0 to disable, 4 MAX) [0]:
```

The recovery configuration for the application watchdog is set using input/output control codes (IOCTLs) that are issued to the `ntwdt` driver.

Watchdog Timer Limitations

The limitations of the watchdog timer mode include:

- In the case of the watchdog timer expiration detected by the system controller, the recovery is attempted only once; there are no further attempts of recovery if the first attempt fails to recover the domain.
- If the application watchdog is enabled and you break into the OpenBoot PROM by issuing the `break` command from the system controller's `sc>` prompt, the system controller automatically disables the watchdog timer.

Note – The system controller displays a console message as a reminder that the watchdog, from the system controller's perspective, is disabled.

However, when you re-enter the Solaris OS, the watchdog timer is still enabled from the Solaris Operating System's perspective. To have both the system controller and the Solaris OS view the same watchdog state, you must use the watchdog application to either enable or disable the watchdog.

- If you perform a dynamic reconfiguration (DR) operation in which a system board containing kernel (permanent) memory is deleted, then you must disable the watchdog timer's application mode before the DR operation and enable it after the DR operation. This is required because Solaris software quiesces all system IO and disables all interrupts during a memory-delete of permanent memory. As a result, system controller firmware and Solaris software can not communicate during the DR operation. Note that this limitation affects neither the dynamic addition of memory nor the deletion of a board not containing permanent memory. In those cases, the watchdog timer's application mode can run concurrently with the DR implementation.

  You can execute the following command to locate the system boards that contain kernel (permanent) memory:

  ```
  sc> cfgadm -lav | grep -i permanent
  ```

- If the Solaris Operating System hangs under the following conditions, the system controller firmware cannot detect the Solaris software hang:
  - Watchdog timer's application mode is set.
  - Watchdog timer is not enabled.
  - No rearming is done by the user.
- The watchdog timer provides partial boot monitoring. You can use the application watchdog to monitor a domain reboot.

  However, domain booting is not monitored for:
  - Bootup after a cold poweron.
  - Recovery of a hung or failed domain.

  In the case of a recovery of a hung or failed domain, a boot failure is not detected and no recovery attempts are made.
- The watchdog timer's application mode provides no monitoring for application startup. In application mode, if the application fails to start up, the failure is not detected and no recovery is provided.

Using the ntwdt Driver

To use the new application watchdog feature, you must install the ntwdt driver. To enable and control the watchdog's application mode, you must program the watchdog system using the LOMIOCDOG*xxx* IOCTLs, described in "Understanding the User API" on page 88.

If the ntwdt driver, as opposed to the system controller, initiates a reset of the Solaris OS on application watchdog expiration, the value of the following property in the ntwdt driver's configuration file (ntwdt.conf) is used:

```
ntwdt-boottimeout="600";
```

In case of a panic, or an expiration of the application watchdog, the ntwdt driver reprograms the watchdog time-out to the value specified in the property.

Assign a value representing a duration that is longer than the time it takes to reboot and perform a crash dump. If the specified value is not large enough, the system controller resets the host if reset is enabled. Note that this reset by the system controller occurs only once.

Understanding the User API

The ntwdt driver provides an application programming interface by using IOCTLs. You must open the /dev/ntwdt device node before issuing the watchdog IOCTLs.

Note – Only a single instance of open() is allowed on /dev/ntwdt. More than one instance of open() will generate the following error message: EAGAIN - The driver is busy, try again.

You can use the following IOCTLs with the watchdog timer:

- LOMIOCDOGTIME
- LOMIOCDOGCTL
- LOMIOCDOGPAT
- LOMIOCDOGSTATE
- LOMIOCALCTL
- LOMIOCALSTATE

Using the Watchdog Timer

Setting the Timeout Period

The LOMIOCDOGTIME IOCTL sets the timeout period of the watchdog. This IOCTL programs the watchdog hardware with the time specified in this IOCTL. You must set the timeout period (LOMIOCDOGTIME) before attempting to enable the watchdog timer (LOMIOCDOGCTL).

The argument is a pointer to an unsigned integer. This integer holds the new timeout period for the watchdog in multiples of 1 second. You can specify any timeout period in the range of 1 second to 180 minutes.

If the watchdog function is enabled, the time-out period is immediately reset so that the new value can take effect. An error (EINVAL) is displayed if the timeout period is less than 1 second or longer than 180 minutes.

Note – The LOMIOCDOGTIME is not intended for general-purpose use. Setting the watchdog time-out to too low a value might cause the system to receive a hardware reset if the watchdog and reset functions are enabled. If the timeout is set too low, the user application must be run with a higher priority (for example, as a real-time thread) and must be rearmed more often to avoid an unintentional expiration.

Enabling or Disabling the Watchdog

The LOMIOCDOGCTL IOCTL enables or disables the watchdog, and it enables or disables the reset capability. See "Finding and Defining Data Structures" on page 90 for the correct values for the watchdog timer.

The argument is a pointer to the lom_dogctl_t structure. This structure is described in greater detail in "Finding and Defining Data Structures" on page 90.

Use the reset_enable member to enable or disable the system reset function. Use the dog_enable member to enable or disable the watchdog function. An error (EINVAL) is displayed if the watchdog is disabled but reset is enabled.

Note – If LOMIOCDOGTIME has not been issued to set up the timeout period prior to this IOCTL, the watchdog is *not* enabled in the hardware.

Rearming the Watchdog

The LOMIOCDOGPAT IOCTL rearms, or pats, the watchdog so that the watchdog starts ticking from the beginning; that is, to the value specified by LOMIOCDOGTIME. This IOCTL requires no arguments. If the watchdog is enabled, this IOCTL must be used at regular intervals that are less than the watchdog timeout, or the watchdog expires.

Getting the State of the Watchdog Timer

The LOMIOCDOGSTATE IOCTL gets the state of the watchdog and reset functions, and retrieves the current time-out period for the watchdog. If LOMIOCDOGSTATE was never issued to set up the timeout period prior to this IOCTL, the watchdog is not enabled in the hardware.

The argument is a pointer to the lom_dogstate_t structure, which is described in greater detail in "Finding and Defining Data Structures" on page 90. The structure members are used to hold the current states of the watchdog reset circuitry and current watchdog timeout period. This timeout period is not the time remaining before the watchdog is triggered.

The LOMIOCDOGSTATE IOCTL requires only that open() be successfully called. This IOCTL can be run any number of times after open() is called, and it does not require any other DOG IOCTLs to have been executed.

Finding and Defining Data Structures

All data structures and IOCTLs are defined in lom_io.h, which is available in the SUNWlomh package.

The data structures for the watchdog timer are shown here:

- The watchdog and reset state data structure is as follows:

CODE EXAMPLE A-1 Watchdog and Reset State Data Structure

```
typedef struct {
        int reset_enable; /* reset enabled if non-zero */
        int dog_enable; /* watchdog enabled if non-zero */
        uint_t dog_timeout; /* Current watchdog timeout */
} lom_dogstate_t;
```

- The watchdog and reset control data structure is as follows:

CODE EXAMPLE A-2 Watchdog and Reset Control Data Structure

```
typedef struct {
        int reset_enable; /* reset enabled if non-zero */
        int dog_enable; /* watchdog enabled if non-zero */
} lom_dogctl_t;
```

Example Watchdog Program

Following is a sample program for the watchdog timer.

CODE EXAMPLE A-3 Example Watchdog Program

```
#include  <sys/types.h>
#include  <fcntl.h>
#include  <unistd.h>
#include  <sys/stat.h>
#include  <lom_io.h>

int main() {
    uint_t timeout = 30; /* 30 seconds */
    lom_dogctl_t dogctl;
    int fd;

    dogctl.reset_enable = 1;
    dogctl.dog_enable = 1;

    fd = open("/dev/ntwdt", O_EXCL);

    /* Set timeout */
    ioctl(fd, LOMIOCDOGTIME, (void *)&timeout);

    /* Enable watchdog */
    ioctl(fd, LOMIOCDOGCTL, (void *)&dogctl);

    /* Keep patting */
    while (1) {
       ioctl(fd, LOMIOCDOGPAT, NULL);
       sleep (5);
    }
    return (0);
}
```

Programming Alarm3

Alarm3 is available to Solaris Operating System users independent of the watchdog mode. Alarm3 or system alarm on and off have been redefined (see TABLE A-1).

Set the value of Alarm3 using the `LOMIOCALCTL` IOCTL. You can program Alarm3 the same way you set and clear Alarm1 and Alarm2.

The following table presents the behavior of Alarm3:

TABLE A-1 Alarm3 Behavior

	Alarm3	Relay	System LED (Green)
Poweroff	On	COM -> NC	Off
Poweron/LOM up	On	COM -> NC	Off
Solaris running	Off	COM -> NO	On
Solaris not running	On	COM -> NC	Off
Host WDT expires	On	COM -> NC	Off
User sets to on	On	COM -> NC	Off
User sets to off	Off	COM -> NO	On

where:

- COM means the common line
- NC means normally closed
- NO means normally open

To summarize the data in the table:

- Alarm3 on = Relay(COM->NC), System LED off
- Alarm3 off = Relay(COM->NO), System LED on

When programmed, you can check Alarm3 or the system alarm with the `showalarm` command and the argument `system`.

For example:

```
sc> showalarm system
system alarm is on
```

The data structure used with the LOMIOCALCTL and LOMIOCALSTATE IOCTLs is as follows:

CODE EXAMPLE A-4 LOMIOCALCTL and LOMIOCALSTATE IOCTL Data Structure

```
#include <fcntl.h>
#include <lom_io.h>

#define LOM_DEVICE "/dev/lom"
#define ALARM_OFF 0
#define ALARM_ON 1

int main() {
    int fd, ret;
    lom_aldata_t ald;
    ald.alarm_no = ALARM_NUM_3;
     ald.state = ALARM_OFF;

     fd = open(LOM_DEVICE, O_RDWR);
     if (fd == -1) {
     printf("Error opening device: %s\n", LOM_DEVICE);
       return (1);
    }

    /* Set Alarm3 to on state */
    ald.state = ALARM_ON;
     ioctl(fd, LOMIOCALCTL, (void *)&ald);

     /* Get Alarm3 state */
    ioctl(fd, LOMIOCALSTATE, (char *)&ald);
     printf("alarm %d state :%d:\n", ald.alarm_no, ald.state);

     /* Set Alarm3 to off state */
     ald.state = ALARM_OFF;
    ioctl(fd, LOMIOCALCTL, (char *)&ald);

     /* Get Alarm3 state */
    ioctl(fd, LOMIOCALSTATE, (char *)&ald);
    printf("alarm %d state :%d:\n", ald.alarm_no, ald.state);

    close (fd);
    return (0);
}
```

Watchdog Timer Error Messages

TABLE A-2 describes watchdog timer error messages that might be displayed and what they mean.

TABLE A-2 Watchdog Timer Error Messages

Error Message	Meaning
`EAGAIN`	An attempt was made to open more than one instance of `open()` on `/dev/ntwdt`.
`EFAULT`	A bad user-space address was specified.
`EINVAL`	A nonexistent control command was requested or invalid parameters were supplied.
`EINTR`	A thread awaiting a component state change was interrupted.
`ENXIO`	The driver is not installed in the system.

APPENDIX **B**

Alarm Relay Output Application Programming Interface

This appendix provides an example program that illustrates how to `get` or `set` the status of the alarms. The application can use the `LOMIOCALSTATE ioctl` function to obtain the status of each alarm and the `LOMIOCALCTL ioctl` function to set the alarms individually. For more details on the alarm indicators, see "Alarm Status Indicators" on page 41.

CODE EXAMPLE B-1 Example Program to `get` and `set` Status of the Alarms

```
#include <sys/types.h>
#include <string.h>
#include <stdlib.h>
#include <sys/unistd.h>
#include <fcntl.h>
#include "lom_io.h"

#define ALARM_INVALID   -1
#define LOM_DEVICE  "/dev/lom"

static void usage();
static void get_alarm(const char *alarm);
static int set_alarm(const char *alarm, const char *alarmval);
static int parse_alarm(const char *alarm);
static int lom_ioctl(int ioc, char *buf);
static char *get_alarmval(int state);
static void get_alarmvals();


main(int argc, char *argv[])
{

        if (argc < 3) {
                usage();
                if (argc == 1)
```

CODE EXAMPLE B-1 Example Program to `get` and `set` Status of the Alarms *(Continued)*

```
                        get_alarmvals();
                exit(1);
        }

        if (strcmp(argv[1], "get") == 0) {
                if (argc != 3) {
                        usage();
                        exit (1);
                }
                        get_alarm(argv[2]);
        }
        else
        if (strcmp(argv[1], "set") == 0) {
                if (argc != 4) {
                        usage();
                        exit (1);
                }
                set_alarm(argv[2], argv[3]);
        } else {
                usage();
                exit (1);
        }
}

static void
usage()
{
        printf("usage: alarm [get|set] [crit|major|minor|user] [on|off]\n");
}

static void
get_alarm(const char *alarm)
{
        ts_aldata_t     ald;
        int altype = parse_alarm(alarm);
        char *val;

        if (altype  == ALARM_INVALID) {
                usage();
                exit (1);
        }

        ald.alarm_no = altype;
        ald.alarm_state = ALARM_OFF;

        lom_ioctl(LOMIOCALSTATE, (char *)&ald);
```

CODE EXAMPLE B-1 Example Program to `get` and `set` Status of the Alarms *(Continued)*

```
        if ((ald.alarm_state != ALARM_OFF) &&
                        (ald.alarm_state != ALARM_ON)) {
                printf("Invalid value returned: %d\n", ald.alarm_state);
                exit(1);
        }

        printf("ALARM.%s = %s\n", alarm, get_alarmval(ald.alarm_state));
}

static int
set_alarm(const char *alarm, const char *alarmstate)
{
        ts_aldata_t     ald;
        int alarmval = ALARM_OFF, altype = parse_alarm(alarm);

        if (altype  == ALARM_INVALID) {
                usage();
                exit (1);
        }

        if (strcmp(alarmstate, "on") == 0)
                alarmval = ALARM_ON;
        else
        if (strcmp(alarmstate, "off") == 0)
                alarmval = ALARM_OFF;
        else {
                usage();
                exit (1);
        }

        ald.alarm_no = altype;
        ald.alarm_state = alarmval;

        if (lom_ioctl(LOMIOCALCTL, (char *)&ald) != 0) {
                printf("Setting ALARM.%s to %s failed\n", alarm, alarmstate);
                return (1);
        } else {
                printf("Setting ALARM.%s successfully set to %s\n", alarm,
alarmstate);
                return (1);
        }
}

static int
parse_alarm(const char *alarm)
{
        int altype;
```

CODE EXAMPLE B-1 Example Program to `get` and `set` Status of the Alarms *(Continued)*

```
        if (strcmp(alarm, "crit") == 0)
                altype = ALARM_CRITICAL;
        else
        if (strcmp(alarm, "major") == 0)
                altype = ALARM_MAJOR;
        else
        if (strcmp(alarm, "minor") == 0)
                altype = ALARM_MINOR;
        else
        if (strcmp(alarm, "user") == 0)
                altype = ALARM_USER;
        else {
                printf("invalid alarm value: %s\n", alarm);
                altype = ALARM_INVALID;
        }

        return (altype);

}

static int
lom_ioctl(int ioc, char *buf)
{
        int fd, ret;

        fd = open(LOM_DEVICE, O_RDWR);

        if (fd == -1) {
                printf("Error opening device: %s\n", LOM_DEVICE);
                exit (1);
        }

        ret = ioctl(fd, ioc, (void *)buf);

        close (fd);

        return (ret);
}

static char *
get_alarmval(int state)
{
        if (state == ALARM_OFF)
                return ("off");
        else
        if (state == ALARM_ON)
```

CODE EXAMPLE B-1 Example Program to `get` and `set` Status of the Alarms *(Continued)*

```
                return ("on");
        else
                return (NULL);
}
static void
get_alarmvals()
{
        get_alarm("crit");
        get_alarm("major");
        get_alarm("minor");
        get_alarm("user");
}
```

Index

I

K

L

M

N

O

P

R

S

T

U

W